NOVEL TECHNOLOGICAL AND METHODOLOGICAL TOOLS FOR THE UNDERSTANDING OF COLLECTIVE BEHAVIORS

EDITED BY: Elio Tuci, Vito Trianni, Simon Garnier and Andrew King

frontiers Research Topics

About Frontiers

Frontiers is more than just an open-access publisher of scholarly articles: it is a pioneering approach to the world of academia, radically improving the way scholarly research is managed. The grand vision of Frontiers is a world where all people have an equal opportunity to seek, share and generate knowledge. Frontiers provides immediate and permanent online open access to all its publications, but this alone is not enough to realize our grand goals.

Frontiers Journal Series

The Frontiers Journal Series is a multi-tier and interdisciplinary set of open-access, online journals, promising a paradigm shift from the current review, selection and dissemination processes in academic publishing. All Frontiers journals are driven by researchers for researchers; therefore, they constitute a service to the scholarly community. At the same time, the Frontiers Journal Series operates on a revolutionary invention, the tiered publishing system, initially addressing specific communities of scholars, and gradually climbing up to broader public understanding, thus serving the interests of the lay society, too.

Dedication to Quality

Each Frontiers article is a landmark of the highest quality, thanks to genuinely collaborative interactions between authors and review editors, who include some of the world's best academicians. Research must be certified by peers before entering a stream of knowledge that may eventually reach the public - and shape society; therefore, Frontiers only applies the most rigorous and unbiased reviews.
Frontiers revolutionizes research publishing by freely delivering the most outstanding research, evaluated with no bias from both the academic and social point of view. By applying the most advanced information technologies, Frontiers is catapulting scholarly publishing into a new generation.

What are Frontiers Research Topics?

Frontiers Research Topics are very popular trademarks of the Frontiers Journals Series: they are collections of at least ten articles, all centered on a particular subject. With their unique mix of varied contributions from Original Research to Review Articles, Frontiers Research Topics unify the most influential researchers, the latest key findings and historical advances in a hot research area! Find out more on how to host your own Frontiers Research Topic or contribute to one as an author by contacting the Frontiers Editorial Office: researchtopics@frontiersin.org

NOVEL TECHNOLOGICAL AND METHODOLOGICAL TOOLS FOR THE UNDERSTANDING OF COLLECTIVE BEHAVIORS

Topic Editors:
Elio Tuci, University of Namur, Belgium
Vito Trianni, Institute of Cognitive Sciences and Technologies, Italian National Research Council, Italy
Simon Garnier, New Jersey Institute of Technology, United States
Andrew King, Swansea University, United Kingdom

Table of Contents

frontiers
in Robotics and AI

Editorial: Novel Technological and Methodological Tools for the Understanding of Collective Behaviors

Elio Tuci[1]*, Vito Trianni[2], Andrew King[3,4] and Simon Garnier[5,6]

[1] Faculty of Informatics, University of Namur, Namur, Belgium, [2] Institute of Cognitive Sciences and Technologies, Italian National Research Council (CNR), Rome, Italy, [3] Department of Biosciences, Swansea University, Swansea, United Kingdom, [4] Department of Biological Sciences, Institute for Communities and Wildlife in Africa, Cape Town, South Africa, [5] Department of Biological Sciences, New Jersey Institute of Technology, Newark, NJ, United States, [6] Department of Biological Sciences, Rutgers University, Newark, NJ, United States

Editorial on the Research Topic

Novel Technological and Methodological Tools for the Understanding of Collective Behaviors

1. INTRODUCTION

The social processes that give rise to coordinated actions of a group of agents and the emergence of global structures—referred to as collective behaviors—are observed in a range of biological and artificial systems. Collective behavior research, therefore, focuses upon a range of different phenomena with the common goal of understanding the dynamics of emergent group level responses, and has resulted in a burgeoning, diverse, and interdisciplinary research community.

Studying collective behaviors in biological and artificial systems is particularly challenging because of their intrinsic complexity, requiring novel approaches that can help unraveling these systems in order to explain how and why certain patterns are produced and maintained. This Research Topic brings together a collection of studies that focus on technological and methodological tools that can support the understanding of collective behaviors. The contributions included within the Research Topic can be broadly categorized as: (i) Review Articles, (ii) Tools and Technologies, and (iii) Empirical Studies.

Our goal is to facilitate the dissemination of ideas, theories, and methods among scientists that share an interest on the study of collective behavior in all its diverse manifestations. It is our hope that, together, this Research Topic and contributions may afford a more complete understanding of the nature of proximate and ultimate causes of collective behaviors in biological systems, and provide opportunity to generate a theoretical framework to engineer robust, resilient, and effective technologies, such as multi-robot systems, smart grids, and sensor networks.

2. REVIEW ARTICLES

Four review articles illustrate the state-of-the-art in the analysis of social dynamics in different research domains.

In Laan et al., the authors review different methodologies to aggregate individual information and restore the collective wisdom when simple averages are not sufficient, explaining when each methodology is applicable to real-world situations. The authors shows that advanced averaging

procedures of the opinions of the members of a large crowd can lead to incredibly accurate collective decisions. However, this accuracy is highly context-dependent and relies on conditions that are often not realistic in practice.

In Moussaïd et al., the authors provide a thorough review of the use and potential for virtual reality and multi-user platforms for new types of experiments in crowd behaviors and describe how these new technologies can transform the way crowd research is conducted. Understanding human crowd dynamics can help urban planners manage crowd safety, ultimately preventing crowd disasters and saving lives.

In Bredeche et al., the authors review research studies focused on a methodology, called embodied evolution, used to design controllers for a group of robots characterized by the use of evolutionary computation techniques in an online fashion. That is, the evaluation, selection, and reproduction cycle runs in a decentralized way on each robot of the group while they are carrying out their task. The authors review a large body of relevant literature and point to a number of open issues that, from their perspective, need to be addressed to further develop this research field into a mature design methodology.

In Tuci et al., the authors provide a comprehensive summary of goals and objectives of the literature on cooperative transport in multi-robot systems. In cooperative transport, a group of robots is required to cooperate in order to transport objects that, due to their mass, shape, or size cannot be transported by single robots. The authors provide an interesting framework to organize a relatively heterogeneous body of work by using the transport strategy as a criterium to classify and sort the research works.

3. TOOLS AND TECHNOLOGIES

Five articles illustrate new methodological tools for the study of collective behaviors.

In Moore et al., the authors illustrate the characteristics of a new library for efficient information-theoretic analysis of collective behaviors. The library proves to be computationally very efficient, inclusive of a larger set of information theoretic measures, and equipped with a suite of wrappers for higher-level programming languages that aim to make it accessible to a wide user-base.

In Boenisch et al., the authors describe a tool to track the movements of bees called BeesBook. Understanding collective behavior of natural systems requires powerful tools to determine the way in which individuals in the collective move and interact with each other. While several tracking softwares are being developed that allow to follow movements and interactions among several animals in a group, few approaches exist for long-term identity-based tracking of individuals. The BeesBook system has been deployed to follow every bee in a colony along a period of several weeks, tracking the movement and interaction of individual insects throughout their whole lifetime.

In Jones et al., the authors illustrate the Xpuck platform to analyse the behavior of a swarm of robots. The ability to synthesize relevant collective behaviors in robot swarms is often bound by the limited computational abilities of robotic platforms,

which do not allow complex information processing (e.g., image analysis for machine vision), or advanced reinforcement learning techniques. The Xpuck platform is therefore an interesting proposal for experimental studies requiring large computational power on the swarm robots, coupling the miniature size of the e-puck platform with the computational power of modern GPUs.

In Nemitz et al., the authors propose a new robotic platform called HoverBot that offers increased functionality while maintaining production costs low. Designed around an innovative locomotion mechanism combining air levitation and magnets, the HoverBot is composed of a single PCB with no mechanical parts making it easy to mass produce and extend with new sensors. This new platform opens up interesting opportunities to implement collective intelligence algorithms on large robotics swarms.

In Bottinelli and Silverberg, the authors describe how methods developed to study granular materials can be applied to difficult-to-analyse patterns of collective motion in biological systems in high density conditions. The authors provide a step-by-step protocol for researchers to create "eigenmodes." These eigenmodes identify hidden long-range motions and localized rearrangements of particles (or any social unit) based solely on their trajectories. This novel approach appears to be a promising new tool for identifying different types of emergent collective motion in biological systems.

4. EMPIRICAL STUDIES

Six articles illustrate results of new experiments focused on collective behavior.

In Hamann, the author study opinion dynamics in a group of mobile robots. A common problem associated with opinion dynamics models when adapted to physical systems—be they natural or artificial—is that the spatial distribution of the agents and their mobility result in spatial correlations that contrast the well-mixed assumptions at the basis of many macroscopic models, making them inappropriate to describe the overall system dynamics. An interesting intuition to grasp the effects of spatial correlation is to include a number of "contrarians" in the population, that is, agents that prefer the minority opinion. With such expedient, the author shows that macroscopic models can be tuned to match the dynamics shown by systems affected by spatial correlations.

For collective behavior studies, a very important ability is the precise identification of leaders in animal groups, as well as the dynamical aspects related to how leaders influence the group movements and how leadership changes from time to time in response to external events. Multiple methods have been proposed in the past, each with its own advantages and drawbacks. In Mwaffo et al., the authors show that a model-free methodology to combine existing methods in a maximum likelihood sense provides an invaluable tool for collective behavior research, allowing to robustly identify leaders from raw positional data.

Object retrieval and gathering has been a hallmark of swarm robotics since its inception in the 1990's. In Strömbom and King,

the authors revive this concept using an algorithm derived from the behavior of sheepdogs. This algorithm uses a feedback loop between a video tracking system and a robot to control the robot movements in relationship to the objects to gather. Results show that the robot can efficiently collect and transport an object to a target location and, more importantly, can adapt its behavior to changing conditions and is robust to noise. This approach to object gathering offers interesting new perspectives for automated swarms of robots.

The use of robots in collective behavior research is a burgeoning area of research. If robots are "accepted" as conspecifics, they allow for experimental manipulations of social interactions. In Bierbach et al., the authors use a bio-mimetic fish in behavioral experiments with surface- and cave-dwelling fish (Poecilia mexicana). They found that both cave- and surface-dwelling fish followed and interacted with the robot when tested in light. However, when tested in darkness, only surface-dwelling fish were attracted to the robot in darkness suggesting the robot fish-replica is providing mostly visual cues. Such work is important because it determines (for this fish system) the mode of feedback between fish and robot, so that the robot (and thus fish) can become controllable by the experimenter.

Nonapeptides (NP) are neurohormones that are known to affect the performance and maintenance of various behaviors in animals, including partner and group preferences. In Ondrasek et al., the authors hypothesized that NP systems may be important mediators to avian collective behaviors, and mapped the distribution of NP receptors in the brain tissue of three flocking bird species—house sparrows, European starlings, and rock doves. The lateral septum, a brain area known to regulate avian flocking was found to have lots of NP receptors in all three species, and in sparrows and starlings the dorsal arcopallium was important too; an area of the brain that we know little about with respect to social behaviors or flocking. Ondrasek et al.'s findings provide an important first step toward the undertaking of neuroecological studies of collective behaviors in birds.

Statistical methods mediated from collective behavior research can be applied also to criminology. The distribution of criminal activities is influenced by the characteristics of the urban environment, but also by indications that a location is (or not) associated with past crimes. In Garnier et al., the authors combine Risk Terrain Modeling, a statistical tool to estimate the relationship between features of the urban environment and crime occurrences, and a model of the spatio-temporal dependence between successive criminal events to predict robberies in a large urban centre. They demonstrate that this two-pronged approach significantly improves upon state-of-the-art methods for predictive policing.

5. CONCLUSIONS

As the reader can clearly notice, the wide range of topics covered in this collection highlights the variety of the research in collective behaviors. The field is largely multi-disciplinary and always crosses the borders of single research domains. As a consequence, it strongly needs new opportunities for gathering together the multiple advances that are constantly proposed—with studies departing from different disciplines—in order to foster cross-fertilization and progress toward a shared understanding of collective behaviors. This Research Topic represents an attempt to provide such a ground: it focused on tools and methods that sometimes have been intended for a specific case (e.g., in the study of animal behavior) but that can be easily generalized to others (e.g., for the design and analysis of robot swarms). This is a pattern that has been followed many times in the past, and we hope that the research work presented here can be of inspiration for further developments in the future.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

frontiers
in Robotics and AI

Analysis of Pairwise Interactions in a Maximum Likelihood Sense to Identify Leaders in a Group

Violet Mwaffo[1], Sachit Butail[2] and Maurizio Porfiri[1]*

[1]Department of Mechanical and Aerospace Engineering, New York University Tandon School of Engineering, Brooklyn, NY, United States, [2]Department of Mechanical Engineering, Northern Illinois University, DeKalb, IL, United States

Collective motion in animal groups manifests itself in the form of highly coordinated maneuvers determined by local interactions among individuals. A particularly critical question in understanding the mechanisms behind such interactions is to detect and classify leader–follower relationships within the group. In the technical literature of coupled dynamical systems, several methods have been proposed to reconstruct interaction networks, including linear correlation analysis, transfer entropy, and event synchronization. While these analyses have been helpful in reconstructing network models from neuroscience to public health, rules on the most appropriate method to use for a specific dataset are lacking. Here, we demonstrate the possibility of detecting leaders in a group from raw positional data in a model-free approach that combines multiple methods in a maximum likelihood sense. We test our framework on synthetic data of groups of self-propelled Vicsek particles, where a single agent acts as a leader and both the size of the interaction region and the level of inherent noise are systematically varied. To assess the feasibility of detecting leaders in real-world applications, we study a synthetic dataset of fish shoaling, generated by using a recent data-driven model for social behavior, and an experimental dataset of pharmacologically treated zebrafish. Not only does our approach offer a robust strategy to detect leaders in synthetic data but it also allows for exploring the role of psychoactive compounds on leader–follower relationships.

Keywords: classification, event synchronization, network, ROC, self-propelled particles, transfer entropy, zebrafish

1. INTRODUCTION

It is generally hypothesized that the movement of animal groups is steered by influential individuals called leaders, which benefit the collective by locating food sources (Giardina, 2008) and protecting against predatory attacks (Partridge, 1982; Ballerini et al., 2008). Further, it is believed that these individuals accomplish these tasks by relying on environmental information available to them rather than social feedback (Dyer et al., 2009; King et al., 2009). Past studies in collective animal behavior have explained the emergence of leadership through several mechanisms, including the availability of extra group knowledge (Krause and Ruxton, 2002; Ioannou et al., 2011), hunger (Krause et al., 1992; Krause, 1993), personality traits (Leblond and Reebs, 2006; Nakayama et al., 2012), and morphophysiological variations (Reebs, 2001).

We work with the definition of leadership by Krause et al. (2000) "as the initiation of new directions of locomotion by one or more individuals which are then readily followed by other group members." Under the assumption that leadership roles within an animal group are consistent

through time and space within the duration of an experimental observation, we seek to identify leaders on the basis of the strength and direction of pairwise interactions among individuals. A leader will be recognized as an individual that exerts a strong one-directional interaction on other group members, while being marginally responsive to their behavior. The interaction between pairs of individuals can be quantified through correlation or information-theoretic measures that capture the directional relationship between the time series of motion data of the individuals. These include cross-correlation (Engel et al., 1990), event synchronization (Quiroga et al., 2002), and information-theoretic measures, such as transfer entropy (Schreiber, 2000), conditional transfer entropy (Sun et al., 2014), maximum entropy (Cavagna et al., 2014), causation entropy (Sun and Bollt, 2014), and union transfer entropy (Anderson et al., 2016).

Each of these measures has its advantages and limitations. Cross-correlation has been successfully used to identify leader–follower relationships from movement data of fish shoals (Krause et al., 2000; Ladu et al., 2014), but it assumes a linear relationship between the time series and is therefore less likely to dissect complex dependencies that consist of varying time delays and non-linear relationships (Ianniello, 1982; Peterson et al., 1998). Event synchronization measures synchronicity between extreme events in the time series (Quiroga et al., 2002) and has been used to identify connectivity structures in atmospheric processes (Malik et al., 2012) and legal policy data (Grabow et al., 2016), under the premise of occurrence of so called extreme events within the time series. Information-theoretic measures, like transfer entropy, have the advantage of being model-free (Steuer et al., 2002; Hlaváčková-Schindler et al., 2007; Vicente et al., 2011), and thereby enable the analysis of time series with varying delays and non-linear relationships. However, since the estimation of these measures requires computing probability distributions, information-theoretic quantities are data hungry (Ito et al., 2011). The duration of observations required to reliably identify relationships between time series increases exponentially with the dimensionality of the dataset (Ito et al., 2011), such that the treatment of multidimensional time series is considerably more challenging than scalar ones.

Animals are likely to communicate within a group through both linear and non-linear dependencies, mediated by unknown delays, making it difficult to pinpoint the specific measure that will perform best for a given dataset of group behavior. Accordingly, all of the above mentioned measures may be useful in identifying leaders at one time or another, and a combined approach that integrates these individual measures could offer a viable approach to study leadership. We detect leader–follower relationships by setting thresholds on average values of pairwise interactions obtained from three different methods: cross-correlation (Engel et al., 1990), event synchronization (Quiroga et al., 2002), and transfer entropy (Schreiber, 2000). To further improve the performance of leader detection beyond any of these methods, we combine them in a maximum likelihood sense to build a single classifier for detecting leaders (Barreno et al., 2008).

Validating this approach would be difficult on real behavioral data, where one may have limited knowledge of, and control on, leadership. Unlike self-propelled particle computer simulations,

where leadership roles can be assigned artificially, identifying leaders within animal groups is hampered by the lack of a ground truth. In this context, we turn to self-propelled particle models to evaluate methods that can identify leaders in group motion. Self-propelled particle models can range from the simplest, where the individuals orient themselves in the general direction of their neighbors (Vicsek et al., 1995; Vicsek and Zafeiris, 2012), to more complex models where interactions include collision avoidance, attraction, and alignment (Aoki, 1982; Couzin et al., 2002, 2005). Data-driven models that incorporate detailed individual dynamics along with species-specific interactions (Gautrais et al., 2009, 2012; Kolpas et al., 2013; Borzí and Wongkaew, 2015; Mwaffo et al., 2015a, 2017; Zienkiewicz et al., 2015a,b; Collignon et al., 2016) provide an even more realistic setup to create such roles and test methods for identifying leaders.

We test our approach on a synthetic dataset comprising simulations of self-propelled particles interacting according to the Vicsek model (Vicsek et al., 1995). A single particle that is not responsive to the rest of the group is assigned the role of a leader. We compare the performance of each classifier as well as the combined classifier in terms of their ability to detect the leader particle. We systematically vary the level of inherent uncertainty and the size of the region of interaction, thereby modulating the degree of coordination within the group (Vicsek et al., 1995). Upon demonstrating the validity of the approach, we investigate its use in the study of realistic data on gregarious fish shoaling. First, we apply the method to detect leaders in an established data-driven model of fish social behavior (Gautrais et al., 2012). Then, we consider experimental data from our group on social behavior of pharmacologically treated zebrafish, in which one fish is exposed to moderate caffeine level to elicit a psychostimulant effect (Fisone et al., 2004; Ferré, 2008). Such a psychostimulant effect could be hypothesized to promote leadership, by potentially reducing social responsiveness and increasing the level of activity of the treated subject, which could be then recognized as a leader by untreated fish (Ladu et al., 2014; Shams and Gerlai, 2016).

The paper is organized as follows. In Section 2, we describe the three classification methods used for studying pairwise interactions in networks of dynamical systems. In Section 3, we explain our approach to detect leadership from raw time series of positional data. We evaluate the performance of all classifiers—individual and combined—on datasets consisting of particles interacting according to the Vicsek model in Section 4. In Section 5, we demonstrate the use of our approach on realistic simulation data and experimental observations on fish collective behavior. We conclude the manuscript with a discussion of the results and performance of the approach.

2. QUANTIFYING PAIRWISE INTERACTIONS IN NETWORKS OF DYNAMICAL SYSTEMS

The process of detecting leaders in a group begins with the measurement of the time series of the individual motion, from which we seek to uncover social interactions. These time series can be

obtained from simulated or experimental data. Specifically, for each individual i, $i = 1, \ldots, N$, where N is the group size, we register a scalar time series $\{x_t^{(i)}\}_{t=1}^{T}$, where T is the duration of the time series and t is the time step. This time series, for example, would represent a salient observable of swimming activity, such as, turn rate, orientation, or positional preference with respect to a target stimulus.
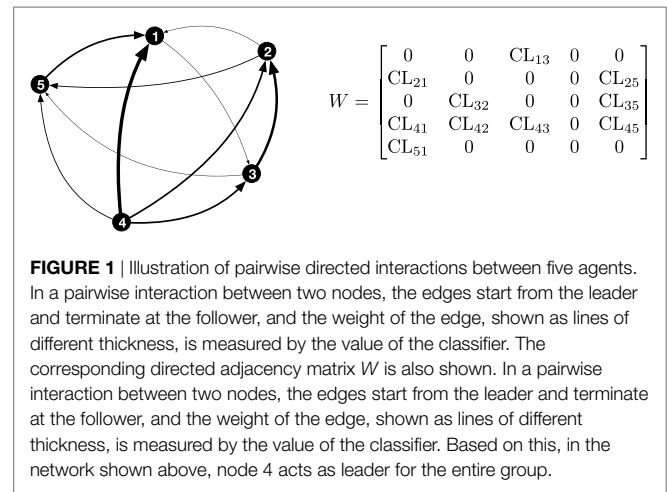
To infer leader–follower relationships between a pair of individuals i and j, we examine three methods, namely, cross-correlation (CC), transfer entropy (TE), and event synchronization (ES). Different from our previous work (Butail et al., 2016), which focused on fish pairs and considered each classification method separately, here we address the more general problem of leader detection in groups in a maximum likelihood sense that integrates the three classifiers. For a pair of individuals and a given method, we construct a one-directional relationship between the individuals, whose magnitude measures the strength of the interaction and whose direction is always from the leader to the follower. In case none of the individuals in the pair is identified as a leader, the strength is set to zero. In general, each method could reveal a different leader–follower relationship for a given pair, and even if methods might agree on who is the leader and who is the follower, the strength of the interaction may vary. We label the strength of the interaction between i and j as $\text{CL}_{ij}^{(\cdot)}$, where the dot specifies the selected method, CC, TE, ES, and CL abbreviates "classifier."

An intuitive representation of leader–follower relationships within the group could be garnered by considering a directed network, where nodes correspond to individuals and weighted directed edges identify the role of each node in the pair (leader versus follower) and the strength of the interaction. As a result, we define the weighted adjacency matrix W, such that $W_{ij}^{(\cdot)} = 0$ if the method detects i as the follower and j as the leader, and $W_{ij}^{(\cdot)} = \text{CL}_{ij}^{(\cdot)} > 0$ if instead i is the leader for the pair ij. The ith row of W has non-zero elements where the pairwise interactions have i as a leader, and the entry corresponds to the value of the classifier. The ith column of W has non-zero elements for the pairwise interactions where i instead is recognized as a follower, and the corresponding entry is the value of the classifier. While it is not possible that both W_{ij} and W_{ji} are non-zero simultaneously, they can both be equal to zero, when the method does not identify a leader in the pair. The weighted adjacency matrix contains all the information that is acquired through the analysis of pairwise interactions, by bookkeeping the role of each node in every possible pairwise interaction and the corresponding strength. **Figure 1** illustrates a network of interaction for a group of five individuals, along with the corresponding weighted adjacency matrix, concisely depicting pairwise leader-follower interactions in the group.

2.1. Cross-Correlation

Cross-correlation measures the similarity between the processes as a function of time delay τ between them (Knapp and Carter, 1976), that is,

$$r_{ij}(\tau) = \frac{\sum_t \left[\left(x_t^{(i)} - \bar{x}^{(i)} \right) \left(x_{t-\tau}^{(j)} - \bar{x}^{(j)} \right) \right]}{\sqrt{\sum_t \left(x_t^{(i)} - \bar{x}^{(i)} \right)^2} \sqrt{\sum_t \left(x_{t-\tau}^{(j)} - \bar{x}^{(j)} \right)^2}}, \quad (1)$$

FIGURE 1 | Illustration of pairwise directed interactions between five agents. In a pairwise interaction between two nodes, the edges start from the leader and terminate at the follower, and the weight of the edge, shown as lines of different thickness, is measured by the value of the classifier. The corresponding directed adjacency matrix W is also shown. In a pairwise interaction between two nodes, the edges start from the leader and terminate at the follower, and the weight of the edge, shown as lines of different thickness, is measured by the value of the classifier. Based on this, in the network shown above, node 4 acts as leader for the entire group.

where $\bar{x}^{(i)}$ and $\bar{x}^{(j)}$ denote the time averages of $x_t^{(i)}$ and $x_t^{(j)}$; the value of t spans the range of overlap between the two time series. The value of delay, τ, that maximizes the cross-correlation $r_{ij}(\tau)$ in equation (1), over a range of values between $-(T-1)$ and $T-1$, is called the time lag between the two time series, that is, $\tau^\star = \text{argmax}_\tau r_{ij}(\tau)$.

When $\tau_{ij}^\star < 0$, we say that $x_t^{(i)}$ anticipates $x_t^{(j)}$, and we identify i as the leader and j as the follower. The numerical value of the corresponding cross-correlation quantifies the strength of the inferred leader–follower interaction, such that, $\text{CL}_{ij}^{\text{CC}} = r_{ij}\left(\tau_{ij}^\star\right)$.

2.2. Transfer Entropy

The computation of transfer entropy requires a probabilistic treatment of the time series. Specifically, we represent each time series $\{x_t^{(i)}\}_{t=1}^{T}$ as a stochastic stationary process $X_t^{(i)}$ taking values in a finite set $\mathcal{X}$. The cardinality of $\mathcal{X}$ is related to the length of the time series, such that longer time series will allow for a high resolution description of the stochastic process, and therefore, a large cardinality. Transfer entropy (Schreiber, 2000) measures the reduction in the uncertainty in predicting one process given the knowledge of another. Transfer entropy from individual j to i is defined as

$$\text{TE}_{j \to i} = \sum_{\mathcal{X}^3} p\left(X_{t+1}^{(i)}, X_t^{(i)}, X_t^{(j)} \right) \log \frac{p\left(X_{t+1}^{(i)} | X_t^{(i)}, X_t^{(j)} \right)}{p\left(X_{t+1}^{(i)} | X_t^{(i)} \right)} \quad (2)$$

Here, $p\left(X_{t+1}^{(i)}, X_t^{(i)}, X_t^{(j)} \right)$ denotes the joint probability of the future and current state of individual i and the current state of individual j; $p\left(X_{t+1}^{(i)} | X_t^{(i)}, X_t^{(j)} \right)$ denotes the conditional probability of the future state of individual i given the current states of both individuals i and j; and $p\left(X_{t+1}^{(i)} | X_t^{(i)} \right)$ denotes the probability of the future state of individual i conditioned on its current state. The probability distributions can be estimated using histograms (Vejmelka and Palus, 2008) or kernel density estimators (Schreiber, 2000). Transfer entropy is a non-negative quantity, which is equal to zero if individual j has no influence on individual i. In this case, $p\left(X_{t+1}^{(i)} | X_t^{(i)} \right) = p\left(X_{t+1}^{(i)} | X_t^{(i)}, X_t^{(j)} \right)$.

We say that i is the leader and j the follower if $\mathrm{TE}_{i\to j} > \mathrm{TE}_{j\to i}$. The value of the, positive, net transfer entropy from the leader to the follower measures the strength of the interaction, that is, $\mathrm{CL}_{ij}^{\mathrm{TE}} = \mathrm{TE}_{i\to j} - \mathrm{TE}_{j\to i}$.

2.3. Extreme-Event Synchronization

Extreme-event synchronization was proposed in Quiroga et al. (2002) to measure synchronicity between signals by comparing the occurrence of extreme events. Briefly, the times when extreme events occur in the two time series for individuals i and j are indexed by $\{t_k^i\}_{k=1}^{m_i}$ and $\{t_k^j\}_{k=1}^{m_j}$, where m_i and m_j are the number of extreme events in the times series of i and j, respectively. These sequences identify the time steps at which the processes exceed a predefined threshold in magnitude; we call such instances extreme events. The number of extreme events for i that occur within a window of duration ξ from those for j are

$$c^{\xi}(i|j) = \sum_{k=1}^{m_i} \sum_{l=1}^{m_j} J_{kl}^{\xi}, \tag{3}$$

where

$$J_{kl}^{\xi} = \begin{cases} 1 & \text{if } 0 < t_l^j - t_k^i \leq \xi, \\ 1/2 & \text{if } t_k^i = t_l^j, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

From the quantity above, we compute event synchronicity and event delay (Quiroga et al., 2002) as follows:

$$Q_{ij}^{\xi} = \frac{c^{\xi}(j|i) + c^{\xi}(i|j)}{\sqrt{m_i m_j}}, \tag{5}$$

$$q_{ij}^{\xi} = \frac{c^{\xi}(j|i) - c^{\xi}(i|j)}{\sqrt{m_i m_j}}. \tag{6}$$

Event synchronicity is symmetric and measures the coupling between individuals i and j; event delay is asymmetric and measures the time lag between extremes events for i and j. By construction, $-1 \leq q_{ij}^{\xi} \leq 1$, such that when $q_{ij}^{\xi} > 0$, the occurrence of extreme events for i systematically precede those for j. We use the sign of event delay to determine leadership, whereby i is the leader if $q_{ij}^{\xi} > 0$. The strength of the interaction is determined by event synchronicity, that is, $\mathrm{CL}_{ij}^{\mathrm{ES}} = Q_{ij}^{\xi}$. By construction, $0 \leq Q_{ij}^{\xi} \leq 1$, with $Q_{ij}^{\xi} = 1$ identifying completely synchronous events.

3. DETECTING LEADERS IN GROUPS

We define group leaders as individuals that on average lead within pairwise interactions with other group members. Using the network representation in **Figure 1**, we identify a group leader as the node with the largest weighted degree, measured as the difference between the weighted out-degree and the weighted in-degree. For node i, the weighted out-degree is the sum of all the pairwise interactions in which the individuals acts as a leader, that is, $\sum_{j=1}^{N} W_{ij}^{(\cdot)}$. The weighted in-degree is the sum of all the pairwise interactions in which the individual acts as a follower, that is, $\sum_{j=1}^{N} W_{ji}^{(\cdot)}$.

As a result, a group leader may not be a leader in every single pairwise interaction, but will have the strongest average effect on the overall group. Specifically, we define the average pairwise interaction for an individual i as

$$\overline{\mathrm{CL}}_i^{(\cdot)} = \frac{1}{N-1} \sum_{j=1}^{N} \left(W_{ij}^{(\cdot)} - W_{ji}^{(\cdot)} \right) \tag{7}$$

and we seek to identify which individual maximizes this quantity. Leaders are classified by setting a threshold $\mathrm{T}_{(\cdot)}$ on the value obtained from equation (7). This combination of average pairwise interaction and the associated threshold constitutes a single classifier.

3.1. Classifier Performance

The performance of a classifier is evaluated in terms of the number of true and false positives and is dependent on the value of the threshold. A visual aid used in comparing different thresholds is the receiver operating characteristic (ROC) curve which plots the number of true positives against false positives for a range of thresholds (Fukunaga, 2013), see, for example, **Figure 2**.

In this respect, a good classifier has few false positives and a large number of true positives for a range of thresholds. Classifier performance can be quantified from the ROC curve by calculating the area under the curve (AUC). A perfect classifier will have 100% true positive rate (TPR) for all values of false positive rate (FPR), and therefore the AUC will be 1. In contrast, a classifier that performs at chance level will have the same number of true and false positives at all combinations and its ROC curve will lie on the diagonal line resulting in an AUC of 0.5.

The optimal threshold value that gives the best performance for a classifier can be estimated from the ROC curve based on several different measures, including distance from the top left corner and the Youden index which maximizes the difference between TPR
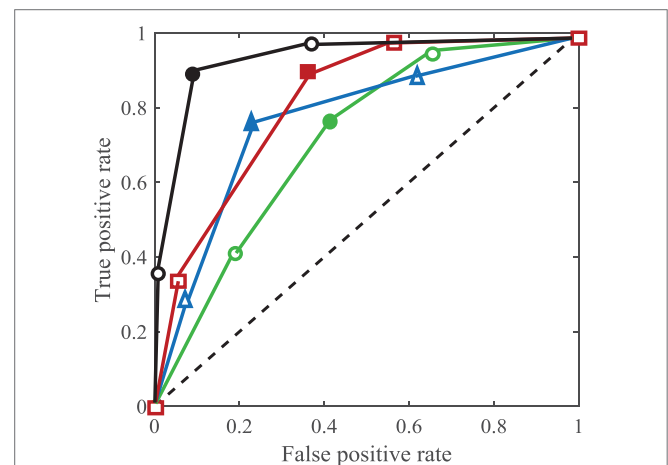
FIGURE 2 | Pictorial illustration of ROC analysis for assessing classifier performance. ROC curves for three hypothetical classifiers are plotted with their respective cutoff points in green, blue, and red. A combined ROC in black is plotted by selecting only three points over the 2^9 produced by the maximum likelihood method. For each curve, the solid marker identifies the operating point, and the empty markers label other cutoff points.

and FPR (Youden, 1950). The corresponding operating point on the ROC curve, which selects the optimal threshold, lies at the maximum vertical distance from the 45° line.

3.2. Combining Classifiers Using Likelihood Ratio

Multiple classifiers can be combined to yield an optimal performance, as illustrated in **Figure 2**, where the black curve is closer to an ideal classifier at the top left corner. Specifically, we combine classifiers in the Neyman–Pearson sense in that the resulting optimal classifier maximizes TPR for a given FPR (Barreno et al., 2008).

The output of a classifier, $\overline{\mathrm{CL}}_i^{(\cdot)}$, and the associated threshold $T_{(\cdot)}$ corresponding to the operating point, can be mapped into the binary choice set $\{0, 1\}$ such that the detection of an individual as a leader corresponds to $\overline{\mathrm{CL}}_i^{(\cdot)} \geq T_{(\cdot)} \equiv 1$ and as a follower to $\overline{\mathrm{CL}}_i^{(\cdot)} < T_{(\cdot)} \equiv 0$. For clarity, we suppress the implicit dependence on the threshold, and denote a classifier simply as $\bar{\mathrm{CL}}_i^{(\cdot)}$. The likelihood ratio for a combination of classifiers $\mathbf{C} = \left(\overline{\mathrm{CL}}^{\mathrm{CC}}, \overline{\mathrm{CL}}^{\mathrm{TE}}, \overline{\mathrm{CL}}^{\mathrm{ES}}\right)$ is defined as $\ell(\mathbf{C}) = P(\mathbf{C}|H_1)/P(\mathbf{C}|H_0)$, where H_1 and H_0 correspond to the hypotheses that the individual being evaluated is a leader or a follower, respectively. In this sense, $P_{(\cdot)}^D = P\left(\overline{\mathrm{CL}}_i^{(\cdot)} = 1|H_1\right)$ corresponds to TPR, and $P_{(\cdot)}^F = P\left(\overline{\mathrm{CL}}_i^{(\cdot)} = 1|H_0\right)$ to FPR. The Neyman–Pearson lemma states that for some value of $\kappa \in (0, \infty)$ and $\gamma \in [0, 1]$, the likelihood ratio test

$$\mathcal{D}(\mathbf{C}) = \begin{cases} 1 & \text{if } \ell(\mathbf{C}) > \kappa, \\ \gamma & \text{if } \ell(\mathbf{C}) = \kappa, \\ 0 & \text{if } \ell(\mathbf{C}) < \kappa \end{cases} \qquad (8)$$

has the highest detection rate, $P(\mathcal{D}(\mathbf{C}) = 1|H_1)$, for a bound on FPR.

The optimal values $\kappa^\star$ and $\gamma^\star$ in the likelihood ratio test are obtained by interpolating between select points on the ROC curve including the operating point, and the (1,1) and (0,0) points on the extreme. These two extreme points identify the cases in which we always classify an individual as a leader, (1,1), or as a follower, (0,0). By interpolating and moving along this new curve, we can tune the false alarm rate. The new ROC curve constructed in this way is called the likelihood-ratio ROC (LR-ROC) (Barreno et al., 2008). Each region of the LR-ROC corresponds to a different decision rule, such that the analyst could locate and use different combinations of classifiers that provide the best performance.

Assuming that the classifiers are conditionally independent, that is $P\left(\overline{\mathrm{CL}}_i^{\mathrm{CC}}, \overline{\mathrm{CL}}_i^{\mathrm{TE}}, \overline{\mathrm{CL}}_i^{\mathrm{ES}}|H_c\right) = P\left(\overline{\mathrm{CL}}_i^{CC}|H_c\right) P\left(\overline{\mathrm{CL}}_i^{\mathrm{TE}}|H_c\right) P\left(\overline{\mathrm{CL}}_i^{\mathrm{ES}}|H_c\right)$, $c \in \{0, 1\}$, we use the true and false positive rates of each to construct the LR-ROC. Specifically, each classifier has two possible outcomes for an individual, that is, an individual can be classified as a follower, when outcome is 0, or leader, when outcome is 1. This results in a total of $2^3 = 8$ possible outcomes for three classifiers. Using the notation $\ell\left(1_{(\cdot)}\right) = P_{(\cdot)}^D/P_{(\cdot)}^F$ to denote the likelihood of classifying an individual as a leader, and $\ell\left(0_{(\cdot)}\right) = \left(1 - P_{(\cdot)}^D\right)/\left(1 - P_{(\cdot)}^F\right)$ to denote the likelihood of

classifying an individual as a follower, we arrange the likelihood ratios in increasing order for eight possible outcomes for three classifiers. From this ordering, for a given value of the false positive rate, we determine the combined true positive rate as the probability maximizing the likelihood ratio, and as such, we construct the combined ROC. The outcomes can be represented with Boolean operators (AND, OR, NOT) to make a combined classifier, where the space of Boolean combinations has cardinality $2^{2^3} = 256$.

In practice, we combine the three classification methods by using three points on their respective ROC. The selection of a small subset of points on the ROC curves is primarily to contain the intensive computational cost associated with searching for the optimal classifier among all possible Boolean combinations (Barreno et al., 2008). Accordingly, we select three points per classifier, close to 25% quartile, 50% quartile and at the operating point of the ROC. Further, in the event that the combined classifier performance measured by the AUC is less than the one of any individual classifier, due to the selection of only three points for the combination, we force the combined method to match the convex hull of the tree classifiers.

Even with three points on each ROC curve, finding the Boolean rule that corresponds to a location on the combined ROC, built using three points per individual classifier,[1] involves searching through a space of $2^{2^9} \approx 1.3 \times 10^{154}$ Boolean combinations of outcomes, which is practically difficult. This does not mean that the combined ROC has no value, since it provides an upper reference bound on which we could test simple Boolean rules that can be easily implemented on a dataset. Such a comparison could be performed by computing the distance between the operating point on the combined ROC and the point that corresponds to a candidate Boolean rule (Khreich et al., 2010).

The maximum likelihood combination of classifiers is a general approach that can accommodate more classifiers, beyond the three considered in this work. However, as the space of Boolean combinations of classifier outcomes rises exponentially (Barreno et al., 2008), the capability of finding the optimal combination becomes practically unfeasible. The combined ROC curve provides an upper bound on which to evaluate candidate Boolean combinations for use in real datasets.

4. CLASSIFYING LEADERS IN VICSEK SELF-PROPELLED PARTICLES

4.1. Modeling Leadership

We adapt the self-propelled particle model proposed by Vicsek (VM) to include leaders, as individuals that do not adjust their heading in response to the rest of the group. Leaders will only change their heading as a function of inherent uncertainty; this behavior could be associated with some prior knowledge of the environment that would manifest into a preference for a given direction. Followers, instead, update their heading based on the response of the group, under the effect of inherent uncertainty. In particular, the model consists of N particles moving in a square of side length L with periodic boundary conditions.

---

[1] Selecting three points per ROC results in 9 binary classifiers to combine, for a total of 2^9 points on the combined ROC.

In the complex plane, the position $\mathbf{x}_i \in \mathbb{C}$ and orientation θ_i of the ith particle changes in time as

$$\mathbf{x}_i(t+1) = \mathbf{x}_i(t) + v e^{I\theta_i(t+1)}, \quad (9a)$$

$$\theta_i(t+1) = \text{Arg}\left[U_i(t)\right] + \eta\zeta, \quad (9b)$$

where $\text{Arg}[\cdot]$ is the phase of a vector; I is the imaginary unit; v is the constant, common speed; $\eta \geq 0$ is the noise intensity; and ζ is uniform random noise in $[-\pi, \pi]$. The vector $U_i(t)$ defines the desired heading of the ith particle, such that

$$U_i(t) = \begin{cases} \frac{1}{|\mathcal{N}_i(t)|} \sum_{j \in \mathcal{N}_i(t)} e^{I\theta_j(t)}, & \text{if } i \text{ is a follower}, \\ e^{I\theta_0}, & \text{if } i \text{ is a leader}, \end{cases} \quad (10)$$

where θ_0 is the preferred heading of the leader. Here, $\mathcal{N}_i(t) = \{j = 1, \dots, N : |\mathbf{x}_i(t) - \mathbf{x}_j(t)| \leq r\}$ is the set of $|\mathcal{N}_i(t)|$ individuals within a circle of radius $r > 0$ from the ith particle. From r and L, one may estimate the average number of neighbors with which a given particle interacts at any time step as $1 + \pi \frac{r^2}{L^2}(N-1)$ (see, for example, Aldana et al., 2007).

Using the VM, we simulate 30 realizations of a group of $N = 5$ self-propelled particles. The simulations are initialized by drawing the particle positions uniformly in a square of length $L = 1$ with their orientations uniformly sampled from $[-\pi, \pi]$. Simulations are performed for 20,000 time steps. Particle turn rate is computed from its heading angle, as $\theta_i(t+1) - \theta_i(t)$ for the ith particle, and utilized to evaluate pairwise interaction using cross-correlation, transfer entropy, and event synchronization. Turn rate is selected

as the key variable for measuring pairwise interactions based on the structure of the VM, in which the only interaction rule is alignment and each particle consistently utilizes its previous heading in the computation of the current heading. As a result, pairwise interactions are likely to manifest in changes of the turn rates.

4.2. Classification

Cross-correlation is computed over the entire length of the time series using the Matlab function *xcorr*. Transfer entropy is computed using PROCESS_NETWORK_v.1.4 software (Ruddell and Kumar, 2009) by estimating the joint probability densities in equation (2) through histograms. The software is run with a total of 18 bins to differentiate the net transfer entropy between group leaders and followers in the VM (see Figure S1 in Supplementary Material). Event synchronization is computed using the MAT-LAB function *Event_sync* developed by Quiroga et al. (2002). To evaluate extreme-event synchronization, similar to Quiroga et al. (2002), the time series of extreme events are extracted from the absolute turn rate, by finding a local maximum over a window of 30 data points. Events between the two time series are considered synchronous if the time lag between them is smaller than half the minimum time lag between successive extreme events in each series (Quiroga et al., 2002). The ROC curves are plotted using the function *perfcurve* available in MATLAB.

Figure 3 illustrates the numerical values of the classification indices in equation (7) for a group of $N = 5$ particles without a leader, with one leader, and with two leaders. For this example, cross-correlation is affected by large standard deviations that may
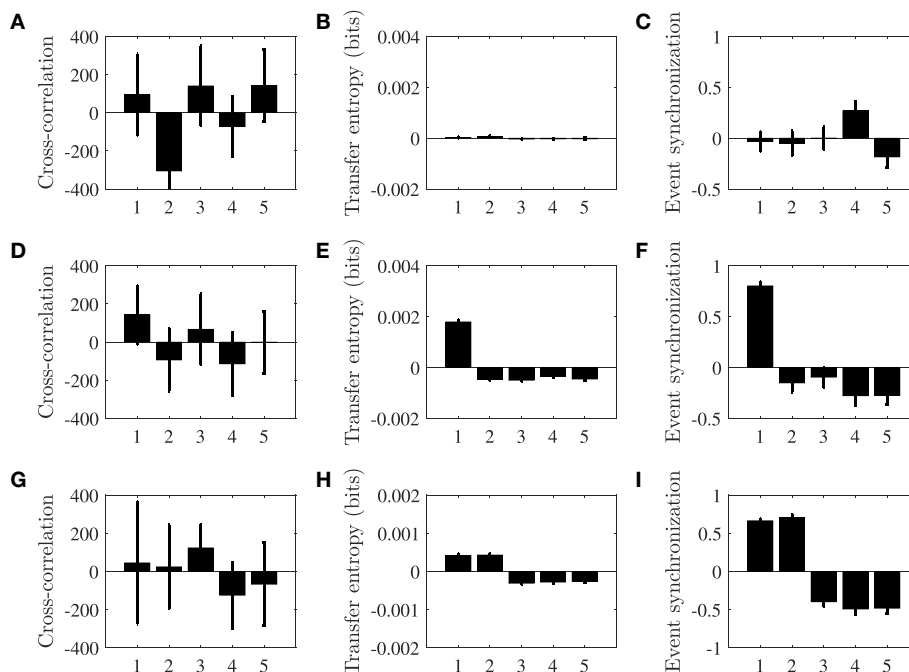
FIGURE 3 | Classification index $\overline{\text{CL}}_i$, for particles $i = 1, \dots, N$ computed for cross-correlation **(A,D,G)**, transfer entropy **(B,E,F)**, event synchronization **(C,F,I)**, without leader **(A–C)**, with one leader **(D–F)**, and with two leaders **(G–I)**. Each simulated group includes five identical particles ($i = 1, \dots, 5$), and the Vicsek model parameters are set to $v = 0.01$, $r = 0.23$, and $\eta = 0.21$. Each bar refers to the mean value of the classifier across 30 simulations, and the error bar is one standard deviation. The numbering of particles that are not leaders is arbitrary, such that in panels **(D–F)** particle 1 is the leader and in panels **(G–I)** particles 1 and 2 are leaders.

mask the success of the detection. Transfer entropy and event synchronization, instead, consistently identify leaders in the group based on the direction and strength of pairwise interactions. To offer some statistical ground for comparing the methods and help assessing the role of model parameters, we next analyze AUC values, focusing on the case of a single leader in the group.

Using ROC, we analyze the performance of the three classification methods in identifying leadership by varying the interaction radius r and the noise intensity η, while keeping the rest of model parameters constant. **Figures 4A–C** present the AUC of the three classifiers as the noise intensity and the radius of interaction are varied. In agreement with our expectations based on the representative case considered in **Figure 3A** cross-correlation is seldom able to correctly identify the leader in the group. For reference, the case displayed in **Figure 3A** has an AUC of 0.51. A likely reason for the limited performance of cross-correlation in detecting leaders in the VM is due to the presence of high-frequency noise in the turn rate, associated with the numerical differentiation of the noise which mediates the orientation update in the model. This noise is likely to suppress linear leader–follower relationships that might be successfully detected using cross-correlation.

Transfer entropy shows excellent performance for every selection of the radius of interaction and a noise intensity between 0.1 and 0.8; for reference the case displayed in **Figure 3B** has an AUC of 1.00. Excessively low noise results into all the particles aligning with the leader's direction in a crystallized formation that does not promote information transfer. In this case, all the particles travel along the constant leader's direction, such that the entropy of each group member is zero. For intensities above 0.8, the particles are nearly independent, such that their orientation update is entirely controlled by noise. In this case, although each particle has a large entropy, the interactions between the particles are masked by individual noise and transfer entropy between any pair of particles vanishes. Increasing the length of the time series could increase the range of noise intensities for which the method can be successful, although dealing with large time series is only realistic for synthetic data. Even if transfer entropy is based on the premise of pairwise interactions, the classification method is successful in isolating the leader for large values of the radius of interaction, which lead to the occurrence of higher-order interactions. This success could be attributed to the use of the average value net transfer entropy across all pairs to construct the classifier, which mitigates the possibility of biases associated with follower-to-follower interactions. Systems composed of a very large number

of particles or the presence of strong heterogeneities could limit the success of the classifier.

Event synchronization demonstrates very good performance for every selection of the radius of interaction and a noise intensity less than 0.4; for reference the case displayed in **Figure 3C** has an AUC of 0.97. For low intensities, noise could manifest in the form of local extreme events in the turn rate which are readily captured by event synchronization. The superior performance of event synchronization with respect to cross-correlation should be attributed to its ability to pick up pairwise leader–follower relationships through varying time delays between extreme events. As noise increases, the frequency of such extreme events becomes too high for establishing faithful relationships between the time series.

The different noise intensity levels at which transfer entropy and event synchronization perform best motivate the need for combining the methods toward a better and more consistent approach to detect leaders in the Vicsek model over more wide range of noise intensities. **Figure 5** demonstrates the performance of the combined method, which yields exact classification for any noise intensity below 0.9.

5. APPLICATIONS TO FISH COLLECTIVE BEHAVIOR

To investigate the applicability of the leader detection approach on fish collective behavior we select two datasets. First, we generate fish-like trajectories from a random walker type model (Gautrais
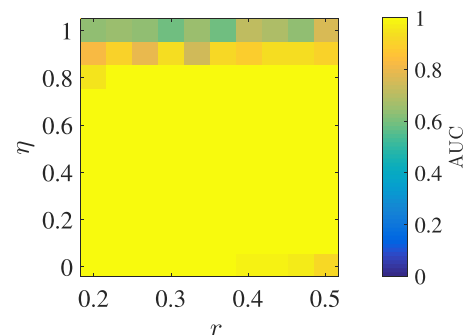
FIGURE 5 | AUC obtained by combining the three classification methods shown in **Figure 4** to detect the single leader, as a function of the radius of interaction r and the noise intensity η, with $N = 5$ and $v = 0.01$.
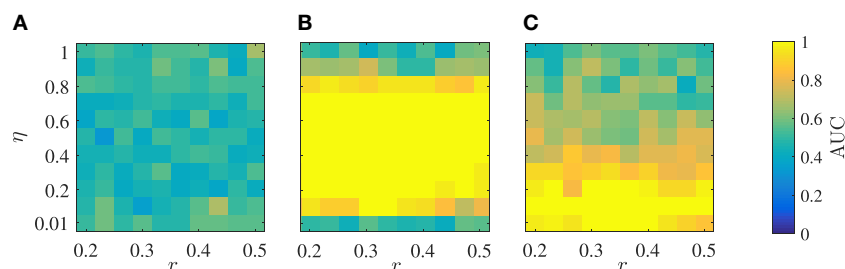
FIGURE 4 | Performance of the three classification methods in detecting the single leader measured by their AUC as a function of the radius of interaction r and the noise intensity η for cross-correlation **(A)**, transfer entropy **(B)**, and event synchronization **(C)**, with $N = 5$ and $v = 0.01$.

et al., 2012) that is able to successfully predict group alignment and average distance in barred flagtails (*Kuhlia mugil*). The data-driven model has five parameters to encapsulate individual swimming, social interactions, and wall interaction. Model parameters are based on selected based on simulations by Gautrais et al. (2012). A single fish is treated as a leader, such that it would not respond to the rest of the group. Second, we utilize trajectories from a group of zebrafish in an experiment where a single fish has been treated with caffeine. In contrast to the trajectories generated using the data-driven model where leadership is systematically controlled, there we explore whether caffeine treatment induces leadership in zebrafish.

5.1. Data-Driven Simulations

The model proposed by Gautrais et al. (2012) offers an authentic data-driven framework to describe the motion of a group of fish. In this model, the turn rate dynamics of a fish is described as a stochastic process modulated by interactions with the environment, which includes members of the group and the tank walls. From the knowledge of the turn rate $\omega_t^{(i)}$ (rad s^{-1}) of fish $i = 1, \ldots, N$, one determines the position $\mathbf{r}^{(i)}$ and orientation $\phi_t^{(i)}$ with respect to a Cartesian coordinate system in $\mathbb{R}^2$ as follows:

$$\frac{d\mathbf{r}_t^{(i)}}{dt} = v \begin{bmatrix} \cos \phi_t^{(i)} \\ \sin \phi_t^{(i)} \end{bmatrix}, \tag{11a}$$

$$\frac{d\phi_t^{(i)}}{dt} = \omega_t^{(i)}, \tag{11b}$$

where v is the common, constant speed.

The instantaneous turn rate at time t is modeled by the mean reverting stochastic differential process (Gautrais et al., 2012; Calovi et al., 2014)

$$d\omega_t^{(i)} = v \left[-\alpha^{(i)} \left(\omega_t^{(i)} - {}^*\omega_t^{(i)} \right) dt + \sigma^{(i)} dW_t^{(i)} \right], \tag{12}$$

where $\alpha^{(i)}$ (s^{-1}) is the rate at which the process returns to its steady state and defines the time scale of the response of a fish to any perturbation; $dW_t^{(i)}$ is the infinitesimal increment of a standard Wiener process resulting in white noise; and $\sigma^{(i)}$ $\left(\text{rad s}^{-3/2}\right)$ is a scaling factor of the Wiener process that measures the level of uncertainty in the motion of a fish. The interaction with the environment is captured by the response function ${}^*\omega_t^{(i)}$ $\left(\text{rad s}^{-1}\right)$

$$
{}^*\omega_t^{(i)} = k_W^{(i)} \frac{\text{sign}\left(\phi_W^{(i)}\right)}{\tau_W^{(i)}} + \frac{1}{N} \sum_{j=1}^N \left[k_v^{(i)} v^{(i)} \sin\left(\phi_t^{(i,j)}\right) \right.
$$
$$
\left. + k_p^{(i)} d_t^{(i,j)} \sin\left(\theta_t^{(i,j)}\right) \right]. \tag{13}
$$

In equation (13), the first term is used to model wall avoidance, and consists of the parameter $k_W^{(i)}$, controlling the intensity of the wall avoidance, $\tau_W^{(i)}$, the time to collision, and $\phi_W^{(i)}$, the angle of incidence with the wall. Both the time to collision and the angle of incidence depend on the instantaneous position and orientation of the fish. The second term in equation (13) measures the

interaction with the rest of the group. Therein, $k_p^{(i)}$ is a parameter controlling the strength of fish attraction toward the group; d_t^{ij} and θ_t^{ij} are the fish interindividual distance and relative angle within the group, respectively; $k_v^{(i)}$ is a parameter controlling the strength of fish alignment with the rest of the group; and $\varphi_t^{ij} = \varphi_t^j - \varphi_t^i$.

We simulate 100 realizations of a group of $N = 5$ fish with a leader. The model is simulated for 120 s using an Euler–Maruyama discretization with time step duration 0.01 s in a circular tank of diameter of 4 m. Orientation is initialized randomly between $[-\pi, \pi)$ and positions are initialized uniformly in the circular domain. The model parameters of the individual turn rate dynamics are taken from Gautrais et al. (2012), that is, $\alpha^{(i)} = 1/0.024 \text{ s}^{-1}$, $\sigma^{(i)} = 28.9 \text{ m}^{-1} \text{ s}^{-1/2}$, and $v = 0.564 \text{ m s}^{-1}$. These values are based on experimental observations on a group of five subjects. We set the first fish as a leader and assign its coupling parameters to zero, that is, $k_p^{(1)} = k_v^{(1)} = 0$, similar to Butail et al. (2016). For the followers, we use $k_p^{(i)} = 0.41 \text{ m}^{-1} \text{ s}^{-1}$, $k_v^{(i)} = 27 \text{ m}^{-1}$, for $i \neq j = 2, \ldots, 5$, to favor coordinated motion, based on results in Zienkiewicz et al. (2015b) and Butail et al. (2016). For all fish, the wall avoidance parameters is set to $k_W^{(i)} = 4.7$, which is larger than the value reported in Gautrais et al. (2012) to reflect the coupling values from Butail et al. (2016). **Figure 6** shows a segment of the trajectories of the simulated group along with the time evolution of their turn rate, which is used for the leader detection process. The computation of the classifiers is analogous to the analysis of the VM, including the number of bins for the computation of transfer entropy that is chosen as 18 (see Figure S2 in Supplementary Material).

In **Figure 7A**, we illustrate the performance of the three classifiers in detecting leadership in the dataset generated using the data-driven model. All the classifiers are successful in detecting a leader beyond chance level, but, as expected from the analysis of the VM, their performance varies. Net transfer entropy and event synchronization, with AUC values at 0.90 and 0.85, respectively, perform better than cross-correlation, with an AUC value of 0.67. **Figure 7B** demonstrates the performance of the combined classifier, generated by selecting three points, indicated in the figure caption, on their respective ROC curve as operating points. Each point on the combined ROC corresponds to a potential combination which can be utilized as a classifier for leadership detection. The combined classifier has an AUC value of 0.99, which is superior than any of the individual classifiers.

In **Table 1**, we show the performance of the best twenty simple Boolean rules with at most three classifiers, ranked based on the distance from the operating point of the combined ROC. For completeness, we display their FPR and TPR. The first five simple Boolean rules have an equivalent performance on this synthetic dataset with an FPR of only 0.08 and a TPR of 0.76.

5.2. Experiments on Pharmacologically Treated Zebrafish

To demonstrate the use of our approach in the study of experimental data on animal behavior, we investigate the possibility that the administration of a psychostimulant compound could elicit leadership in a group of fish. Specifically, we consider experimental
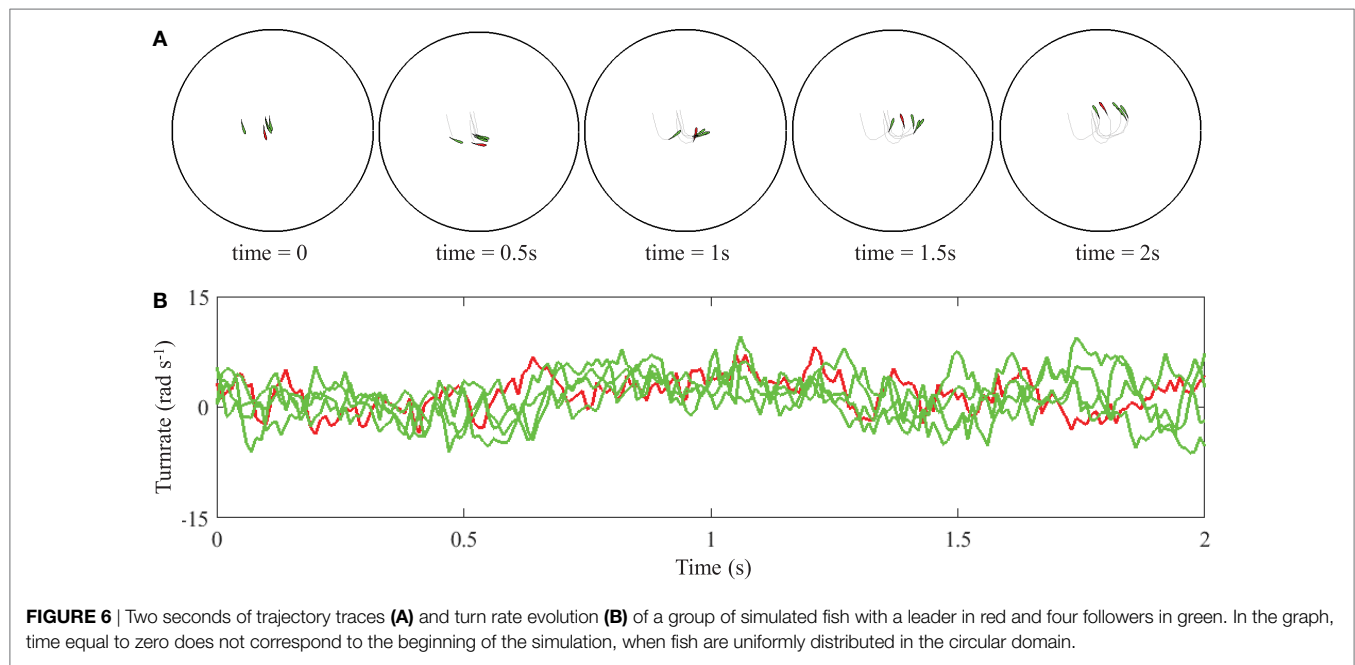
FIGURE 6 | Two seconds of trajectory traces **(A)** and turn rate evolution **(B)** of a group of simulated fish with a leader in red and four followers in green. In the graph, time equal to zero does not correspond to the beginning of the simulation, when fish are uniformly distributed in the circular domain.
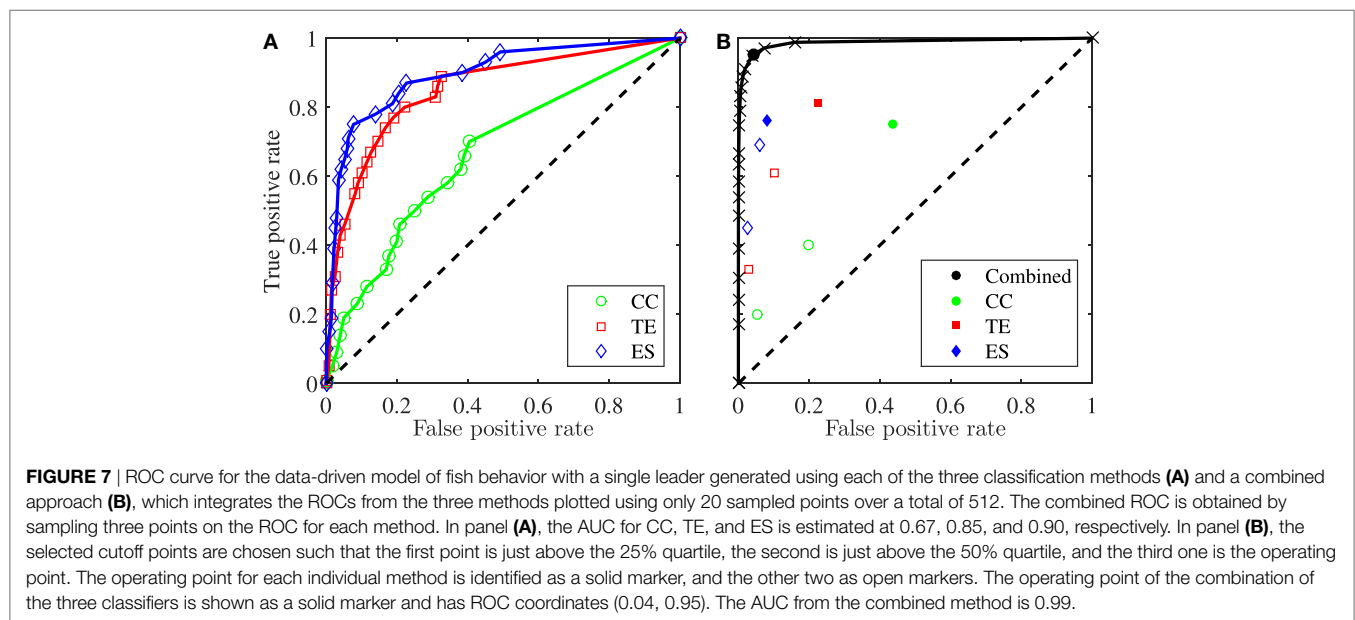
FIGURE 7 | ROC curve for the data-driven model of fish behavior with a single leader generated using each of the three classification methods **(A)** and a combined approach **(B)**, which integrates the ROCs from the three methods plotted using only 20 sampled points over a total of 512. The combined ROC is obtained by sampling three points on the ROC for each method. In panel **(A)**, the AUC for CC, TE, and ES is estimated at 0.67, 0.85, and 0.90, respectively. In panel **(B)**, the selected cutoff points are chosen such that the first point is just above the 25% quartile, the second is just above the 50% quartile, and the third one is the operating point. The operating point for each individual method is identified as a solid marker, and the other two as open markers. The operating point of the combination of the three classifiers is shown as a solid marker and has ROC coordinates (0.04, 0.95). The AUC from the combined method is 0.99.

data by our group (submitted work—data available upon request) on the collective behavior of caffeine-treated zebrafish swimming in a shallow water circular tank. The experimental procedure was carried out under protocol number 13-1424, approved by the University Animal Welfare Committee (UAWC) of New York University. In the literature, a number of studies have explored the effects of this psychoactive compound on the individual behavior of this popular animal model, but the effect of caffeine on zebrafish social behavior has yet to be fully understood (García-Pardo et al., 2015).

In our experiment, we test 10 groups of five fish, in which only one of the subjects is treated with caffeine at 25 mg/l concentration level. Fish motion is recorded from an overhead view at 40 frames per second for 5 min of experiments. A Daubechies wavelet filter is first applied to the fish centroid positions, and the turn rate of each fish, ω_t^i with $i = 1, \ldots, 5$, is consequently estimated from the curvature of the trajectory (Mwaffo et al., 2015b). Following (Butail et al., 2016), data are down-sampled to a sampling period of 0.2 s to minimize the effect of measurement noise on the interactions. The number of bin is set at 18 to ensure consistency with respect to the simulation results presented earlier.

To implement our method on experimental data of fish treated with caffeine, we select the Boolean rule $\neg$ CC $\wedge$ TE $\vee$ ES in **Table 1**. This selection is based on the following reasons: (i) this Boolean rule shows the best performance on the synthetic data generated by the data-driven model of fish social behavior, as

TABLE 1 | Performance of 20 select Boolean rules on the synthetic dataset of data-driven model of fish social behavior.

Rank	Boolean rule	FPR	TPR	Distance
1	ES	0.08	0.76	0.19
2	CC ∧ ES	0.08	0.76	0.19
3	¬ CC ∨ ES	0.08	0.76	0.19
4	¬ CC ∧ TE ∨ ES	0.08	0.76	0.19
5	¬ CC ∧ ¬ TE ∨ ES	0.08	0.76	0.19
6	TE	0.23	0.81	0.23
7	TE ∧ CC	0.23	0.81	0.23
8	TE ∨ ¬ CC	0.23	0.81	0.23
9	TE ∨ ES	0.28	0.96	0.23
10	¬ CC ∨ TE ∨ ES	0.28	0.96	0.23
11	CC ∧ TE ∨ ES	0.28	0.96	0.23
12	TE ∧ ES	0.03	0.61	0.33
13	CC ∧ TE ∧ ES	0.03	0.61	0.33
14	¬ CC ∨ TE ∧ ES	0.03	0.61	0.33
15	CC	0.44	0.75	0.44
16	TE ∧ ¬ ES	0.20	0.20	0.76
17	CC ∧ TE ∧ ¬ ES	0.20	0.20	0.76
18	¬ CC ∨ TE ∧ ¬ ES	0.20	0.20	0.76
19	¬ TE ∨ ES	0.81	0.80	0.78
20	¬ CC ∨ ¬ TE ∨ ES	0.81	0.80	0.78

The Boolean rules consist of logical combinations of individual classifiers corresponding to their operating points. Performance of each Boolean rule is ranked with respect to its distance from the operating point of the combined ROC, such that, larger distance means a worse classifier.

shown in **Table 1** and (ii) it combines TE and ES, which are found to complement each other in the classification of leaders and followers in the VM for the entire parameter space, as shown in **Figure 4**. Although the other four best rules in **Table 1** have the same performance on the simulated dataset, they do not use TE, which is important for detecting leaders in instances of the VM characterized by limited coordination between the particles. The thresholds of CC, TE, and ES used to implement the Boolean rule on experimental data are obtained from the ROCs for the synthetic data generated by the data-driven model of fish social behavior. Specifically, we scale the operating points on those ROCs by the maximum values of CC, TE, and ES in the simulation and apply these thresholds to experimental data, which is also scaled by their corresponding maximum values.

In Table S1 in Supplementary Material, we summarize the results of the combined detection rule. For 10 out of the 10 experiments, we find that the Boolean rule ¬ CC ∧ TE ∨ ES identifies the caffeine-treated fish as a leader for the group. By comparing the fraction of experiments in which the treated fish is identified as a leader (10/10) with chance (1/5) using a t-test, we cannot dismiss the hypothesis that caffeine treatment is a determinant of leadership ($t(9) = 1$, $p < 0.01$). This result could be explained by the psychostimulant effects of caffeine, which, similar to other psychoactive compounds, like lysergic acid diethylamide and 3,4-methylenedioxymethamphetamine, might modulate social responsiveness (Shams and Gerlai, 2016). Also, we may propose that caffeine could enhance fish activity and produce an increase in the frequency of fast and sudden turning maneuvers (Wong et al., 2010; Gupta et al., 2014). It is possible that the hyperactivity of the treated fish could be perceived by untreated fish as an indicator of fitness, boldness, or high social status, thereby favoring its appraisal as a group leader (Ladu et al., 2014).

6. CONCLUSION

Here, we investigate the possibility of detecting leaders in animal groups from raw position data of each individual. Our approach to leadership detection builds on the measurement of pairwise interactions between each pair of individuals to isolate individuals that exert maximum net influence over the rest of the group based on a receiving operating curve. Pairwise interactions are quantified using three independent methods—cross-correlation, transfer entropy, and event synchronization—that are cogently integrated to maximize our success to identify leaders from raw data. In the technical literature, each of these methods has been found to have differential success in the study of connectivity patterns: we hypothesize that their combination in a maximum likelihood sense would help bring to light their specific advantages and mitigate their limitations.

We demonstrate our approach through the systematic study of self-propelled particles described using the classical Vicsek model (Vicsek et al., 1995), in which particles update their orientation as a function of their neighbors and additive noise. The leader is modeled as a particle that has additional knowledge about a specific direction to take, thereby maintaining its orientation, irrespective of the rest of the group. We systematically elucidate the role of the radius of interaction and the noise intensity on the success of each of the three methods to detect the leader. While cross-correlation typically fails to accurately identify the leader, the combination of transfer entropy and event synchronization demonstrates excellent performance for any parameter selection. From raw time series, we show the possibility of exactly detecting a leader from small to large noise intensities, encapsulating disordered and ordered patterns, and form small to large radii of interactions, describing sparse to fully connected networks of followers. The possibility of successfully detecting a single leader is not masked by introducing mild heterogeneities in the groups.[2]

Based on the success of our combined approach, we tackle two realistic datasets of fish social behavior. First, we demonstrate the ability to detect a leader in a synthetic dataset generated using a data-driven model (Gautrais et al., 2012; Calovi et al., 2014), in which the turn rate of each fish is described as a mean reverting diffusion process. Through our combined approach, we are successful in precisely isolating the leader from the rest of the group. Next, we study an experimental dataset on pharmacologically treated fish, in which one of the subjects is administered caffeine to elicit a psychostimulant effect that could enhance activity and trigger leadership. In agreement with the premise of the experiment, through the application of our combined approach, we find that caffeine-treated subjects are more likely to emerge as leaders of the group.

Our approach of identifying leaders *via* the strength of interactions over experimental time assumes that the leaders are consistent throughout the entire observation, in time and in space, which may not be always the case (Nakayama et al., 2012). When these

---

[2]We tested our approach with a group of 5 simulated fish whose parameters were chosen within ±10% of their nominal values used to generate **Figure 7**. Our results show similar performance for each classifier as well as the improvement in performance from the combined classifier—see Figure S3 in Supplementary Material.

conditions lose validity, one may seek to partition the observation into contiguous measurements and implement the approach separately, on each measurement. If data are available at high resolution, the analysis should reveal how leadership varies in the group during the observation.

Another important assumption of our approach is that a group member can either be a leader or a follower, which may not always be the case (Rosenthal et al., 2015). Although it is possible to mark an interaction as leaderless based on the value of the interaction strength, computing the baseline for such values may require experiments that tie leadership with other personality based traits. Understanding the number of leaders that the method can detect is also an area that requires further research. While our method is able to identify single leaders in small and large groups,[3] its applicability to the study of groups with multiple leaders may pose some technical challenges due to the possibility of large correlation lengths and groups splits (DeLellis et al., 2013).

Further, leaders in our simulated datasets assume a singular role in the group, whereby they are not influenced by the rest of the individuals. A scenario may exist where leaders could act on information provided by a subset of neighbors, designated as informed followers, in the absence of consensus (Cucker and Huepe, 2008). It is likely that in such scenarios, the interaction strength will be lowered as compared to the directed relationships simulated here, thereby challenging the process of inference based on ROC curves.

This study significantly strengthens our methodological toolbox to study leadership in animal groups, by empowering analysts with a model-free framework to investigate the basis and determinants of leadership. This effort significantly expands on our previous work (Butail et al., 2016), which is limited to pairs and

---

[3] We evaluated our approach with a group of 20 simulated fish, which shows similar performance for each classifier as well as the improvement in performance from the combined classifier—see Figure S4 in Supplementary Material.

REFERENCES

Aldana, M., Dossetti, V., Huepe, C., Kenkre, V., and Larralde, H. (2007). Phase transitions in systems of self-propelled agents and related network models. *Phys. Rev. Lett.* 98, 095702. doi:10.1103/PhysRevLett.98.095702

Anderson, R. P., Jimenez, G., Bae, J. Y., Silver, D., Macinko, J., and Porfiri, M. (2016). Understanding policy diffusion in the US: an information-theoretical approach to unveil connectivity structures in slowly evolving complex systems. *SIAM J. Appl. Dyn. Syst.* 15, 1384–1409. doi:10.1137/15M1041584

Aoki, I. (1982). A simulation study on the schooling mechanism in fish. *Bull. Jpn. Soc. Sci. Fish.* 48, 1081–1088. doi:10.2331/suisan.48.1081

Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., et al. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: evidence from a field study. *Proc. Natl. Acad. Sci. U.S.A.* 105, 1232–1237. doi:10.1073/pnas.0711437105

Barreno, M., Cardenas, A., and Tygar, J. D. (2008). "Optimal ROC curve for a combination of classifiers," in *Proceedings of the 22nd Annual Conference on Neural Information Processing Systems, December 2008*, Vancouver, British Columbia, Canada, 57–64.

Borzí, A., and Wongkaew, S. (2015). Modeling and control through leadership of a refined flocking system. *Math. Models Methods Appl. Sci.*, 25, 255–282.

Butail, S., Mwaffo, V., and Porfiri, M. (2016). Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Phys. Rev. E* 93, 042411. doi:10.1103/PhysRevE.93.042411

Calovi, D. S., Lopez, U., Ngo, S., Sire, C., Chaté, H., and Theraulaz, G. (2014). Swarming, schooling, milling: phase diagram of a data-driven fish school model. *New J. Phys.* 16, 015026. doi:10.1088/1367-2630/16/1/015026

Cavagna, A., Giardina, I., Ginelli, F., Mora, T., Piovani, D., Tavarone, R., et al. (2014). Dynamical maximum entropy approach to flocking. *Phys. Rev. E* 89, 042707. doi:10.1103/PhysRevE.89.042707

Collignon, B., Séguret, A., and Halloy, J. (2016). A stochastic vision-based model inspired by zebrafish collective behaviour in heterogeneous environments. *R. Soc. Open Sci.* 3, 150473. doi:10.1098/rsos.150473

Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature* 433, 513–516. doi:10.1038/nature03236

Couzin, I. D., Krause, J., James, R., Ruxton, G. D., and Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *J. Theor. Biol.* 218, 1–11. doi:10.1006/jtbi.2002.3065

Cucker, F., and Huepe, C. (2008). Flocking with informed agents. *Math. Action* 218, 1–25. doi:10.5802/msia.1

DeLellis, P., Porfiri, M., and Bollt, E. M. (2013). Topological analysis of group fragmentation in multiagent systems. *Phys. Rev. E* 87, 022818. doi:10.1103/PhysRevE.87.022818

Dyer, J. R., Johansson, A., Helbing, D., Couzin, I. D., and Krause, J. (2009). Leadership, consensus decision making and collective behaviour in humans. *Anim. Behav.* 364, 781–789. doi:10.1098/rstb.2008.0233

Engel, A. K., König, P., Gray, C. M., and Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: inter-columnar interaction as determined

ETHICS STATEMENT

AUTHOR CONTRIBUTIONS

All the authors designed the study, performed the analysis of the data, and wrote the manuscript.

FUNDING

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at http://journal.frontiersin.org/article/10.3389/frobt.2017.00035/full#supplementary-material.

by cross-correlation analysis. *Eur. J. Neurosci.* 2, 588–606. doi:10.1111/j.1460-9568.1990.tb00449.x

Ferré, S. (2008). An update on the mechanisms of the psychostimulant effects of caffeine. *J. Neurochem.* 105, 1067–1079. doi:10.1111/j.1471-4159.2007.05196.x

Fish, F. E., Fegely, J. F., and Xanthopoulos, C. J. (1991). Burst-and-coast swimming in schooling fish (*Notemigonus crysoleucas*) with implications for energy economy. *Comp. Biochem. Physiol. A Physiol.* 100, 633–637. doi:10.1016/0300-9629(91)90382-M

Fisone, G., Borgkvist, A., and Usiello, A. (2004). Caffeine as a psychomotor stimulant: mechanism of action. *Cell. Mol. Life Sci.* 61, 857–872. doi:10.1007/s00018-003-3269-3

Fukunaga, K. (2013). *Introduction to Statistical Pattern Recognition*. San Diego: Academic Press.

García-Pardo, M. P., Blanco-Gandía, M. C., Valiente-Lluch, M., Rodríguez-Arias, M., Miñarro, J., and Aguilar, M. A. (2015). Long-term effects of repeated social stress on the conditioned place preference induced by MDMA in mice. *Prog. Neuropsychopharmacol. Biol. Psychiatry* 63, 98–109. doi:10.1016/j.pnpbp.2015.06.004

Gautrais, J., Ginelli, F., Fournier, R., Blanco, S., Soria, M., Chaté, H., et al. (2012). Deciphering interactions in moving animal groups. *PLoS Comput. Biol.* 8:e1002678. doi:10.1371/journal.pcbi.1002678

Gautrais, J., Jost, C., Soria, M., Campo, A., Motsch, S., Fournier, R., et al. (2009). Analyzing fish movement as a persistent turning walker. *J. Math. Biol.* 58, 429–445. doi:10.1007/s00285-008-0198-7

Giardina, I. (2008). Collective behavior in animal groups: theoretical models and empirical studies. *HFSP J.* 2, 205–219. doi:10.2976/1.2961038

Grabow, C., Macinko, J., Silver, D., and Porfiri, M. (2016). Detecting causality in policy diffusion processes. *Chaos* 26, 083115. doi:10.1063/1.4961067

Gupta, P., Khobragade, S. B., Shingatgeri, V. M., Rajaram, S. M. (2014). Assessment of locomotion behavior in adult zebrafish after acute exposure to different pharmacological reference compounds. *Drug Dev. Ther.* 5, 127–133. doi:10.4103/2394-2002.139626

Hlaváčková-Schindler, K., Paluš, M., Vejmelka, M., and Bhattacharya, J. (2007). Causality detection based on information-theoretic approaches in time series analysis. *Phys. Rep.* 441, 1–46. doi:10.1016/j.physrep.2006.12.004

Ianniello, J. (1982). Time delay estimation via cross-correlation in the presence of large estimation errors. *IEEE Trans. Acoust.* 30, 998–1003. doi:10.1109/TASSP.1982.1163992

Ioannou, C. C., Couzin, I. D., James, R., Croft, D. P., and Krause, J. (2011). Social organisation and information transfer in schooling fish. *Fish Cogn. Behav.* 217–239. doi:10.1002/9781444342536.ch10

Ito, S., Hansen, M. E., Heiland, R., Lumsdaine, A. M., Litke, A., and Beggs, J. M. (2011). Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. *PLoS ONE* 6:e27431. doi:10.1371/journal.pone.0027431

Khreich, W., Granger, E., Miri, A., and Sabourin, R. (2010). Iterative Boolean combination of classifiers in the ROC space: an application to anomaly detection with HMMs. *Pattern Recognit.* 43, 2732–2752. doi:10.1016/j.patcog.2010.03.006

King, A. J., Johnson, D. D. P., and Van Vugt, M. (2009). The origins and evolution of leadership. *Curr. Biol.* 19, R911–R916. doi:10.1016/j.cub.2009.07.027

Knapp, C., and Carter, G. (1976). The generalized correlation method for estimation of time delay. *IEEE Trans. Acoust.* 24, 320–327. doi:10.1109/TASSP.1976.1162830

Kolpas, A., Busch, M., Li, H., Couzin, I. D., Petzold, L., and Moehlis, J. (2013). How the spatial position of individuals affects their influence on swarms: a numerical comparison of two popular swarm dynamics models. *PLoS ONE* 8:e58525. doi:10.1371/journal.pone.0058525

Krause, J. (1993). The influence of hunger on shoal size choice by three-spined sticklebacks, *Gasterosteus aculeatus*. *J. Fish Biol.* 43, 775–780. doi:10.1111/j.1095-8649.1993.tb01154.x

Krause, J., Bumann, D., and Todt, D. (1992). Relationship between the position preference and nutritional state of individuals in schools of juvenile roach (*Rutilus rutilus*). *Behav. Ecol. Sociobiol.* 30, 177–180. doi:10.1007/BF00166700

Krause, J., Hoare, D., Krause, S., Hemelrijk, C. K., and Rubenstein, D. I. (2000). Leadership in fish shoals. *Fish Fish.* 1, 82–89. doi:10.1111/j.1467-2979.2000.tb00001.x

Krause, J., and Ruxton, G. D. (2002). *Living in Groups*. Oxford: Oxford University Press.

Ladu, F., Butail, S., Macrí, S., and Porfiri, M. (2014). Sociality modulates the effects of ethanol in zebra fish. *Alcohol. Clin. Exp. Res.* 38, 1–9. doi:10.1111/acer.12432

Leblond, C., and Reebs, S. G. (2006). Individual leadership and boldness in shoals of golden shiners (*Notemigonus crysoleucas*). *Behaviour* 143, 1263–1280. doi:10.1163/156853906778691603

Malik, N., Bookhagen, B., Marwan, N., and Kurths, J. (2012). Analysis of spatial and temporal extreme monsoonal rainfall over South Asia using complex networks. *Clim. Dyn.* 39, 971–987. doi:10.1007/s00382-011-1156-4

Mwaffo, V., Anderson, R. P., Butail, S., and Porfiri, M. (2015a). A jump persistent turning walker to model zebrafish locomotion. *J. R. Soc. Interface* 12, 20140884. doi:10.1098/rsif.2014.0884

Mwaffo, V., Butail, S., Di Bernardo, M., and Porfiri, M. (2015b). Measuring zebrafish turning rate. *Zebrafish* 12, 250–254. doi:10.1089/zeb.2015.1081

Mwaffo, V., Butail, S., and Porfiri, M. (2017). In-silico experiments of zebrafish behaviour: modeling swimming in three dimensions. *Sci. Rep.* 7, 39877. doi:10.1038/srep39877

Nakayama, S., Harcourt, J. L., Johnstone, R. A., and Manica, A. (2012). Initiative, personality and leadership in pairs of foraging fish. *PLoS ONE* 7:e36606. doi:10.1371/journal.pone.0036606

Partridge, B. L. (1982). The structure and function of fish schools. *Sci. Am.* 246, 114–123. doi:10.1038/scientificamerican0682-114

Peterson, B. M., Wanders, I., Horne, K., Collier, S., Alexander, T., Kaspi, S., et al. (1998). On uncertainties in cross-correlation lags and the reality of wavelength-dependent continuum lags in active galactic nuclei. *Publ. Astron. Soc. Pac.* 110, 660. doi:10.1086/316177

Quiroga, R. Q., Kreuz, T., and Grassberger, P. (2002). Event synchronization: a simple and fast method to measure synchronicity and time delay patterns. *Phys. Rev. E* 66, 041904. doi:10.1103/PhysRevE.66.041904

Reebs, S. G. (2001). Influence of body size on leadership in shoals of golden shiners, *Notemigonus crysoleucas*. *Behaviour* 138, 797–809. doi:10.1163/156853901753172656

Rosenthal, S. B., Twomey, C. R., Hartnett, A. T., Wu, H. S., and Couzin, I. D. (2015). Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion. *Proc. Natl. Acad. Sci. U.S.A.* 112, 4690–4695. doi:10.1073/pnas.1420068112

Ruddell, B. L., and Kumar, P. (2009). Ecohydrologic process networks: 1. Identification. *Water Resour. Res.* 45, W03419. doi:10.1029/2008WR007279

Schreiber, T. (2000). Measuring information transfer. *Phys. Rev. Lett.* 85, 461. doi:10.1103/PhysRevLett.85.461

Shams, S., and Gerlai, R. (2016). Pharmacological manipulation of shoaling behavior in zebrafish. *Curr. Psychopharmacol.* 5, 180–193. doi:10.2174/2211556005666160607094906

Steuer, R., Kurths, J., Daub, C. O., Weise, J., and Selbig, J. (2002). The mutual information: detecting and evaluating dependencies between variables. *Bioinformatics* 18, S231–S240. doi:10.1093/bioinformatics/18.suppl_2.S231

Sun, J., and Bollt, E. M. (2014). Causation entropy identifies indirect influences, dominance of neighbors and anticipatory couplings. *Physica D* 267, 49–57. doi:10.1016/j.physd.2013.07.001

Sun, Y., Rossi, L. F., Shen, C. C., Miller, J., Wang, X. R., Lizier, J. T., et al. (2014). *Information Transfer in Swarms with Leaders*. arXiv Preprint arXiv:1407.0007.

Vejmelka, M., and Palus, M. (2008). Inferring the directionality of coupling with conditional mutual information. *Phys. Rev. E* 77, 026214. doi:10.1103/PhysRevE.77.026214

Vicente, R., Wibral, M., Lindner, M., and Pipa, G. (2011). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *J. Comput. Neurosci.* 30, 45–67. doi:10.1007/s10827-010-0262-3

Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., and Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.* 75, 1226–1229. doi:10.1103/PhysRevLett.75.1226

Vicsek, T., and Zafeiris, A. (2012). Collective motion. *Phys. Rep.* 517, 71–140. doi:10.1016/j.physrep.2012.03.004

Wong, K., Elegante, M., Bartels, B., Elkhayat, S., Tien, D., Roy, S., et al. (2010). Analyzing habituation responses to novelty in zebrafish (*Danio rerio*). *Behav. Brain Res.* 208, 450–457. doi:10.1016/j.bbr.2009.12.023

Youden, W. J. (1950). Index for rating diagnostic tests. *Cancer* 3, 32–35. doi:10.1002/1097-0142(1950)3:1<32::AID-CNCR2820030106>3.0. CO;2-3

Zienkiewicz, A., Barton, D. A. W., Porfiri, M., and di Bernardo, M. (2015a). Data-driven stochastic modelling of zebrafish locomotion. *J. Math. Biol.* 71, 1081–1105. doi:10.1007/s00285-014-0843-2

Zienkiewicz, A., Barton, D. A. W., Porfiri, M., and di Bernardo, M. (2015b). Leadership emergence in a data-driven model of zebrafish shoals with speed modulation. *Eur. Phys. J. Special Top.* 224, 3343–3360. doi:10.1140/epjst/e2015-50093-5

HoverBots: Precise Locomotion Using Robots That Are Designed for Manufacturability

Markus P. Nemitz[1,2]*, Mohammed E. Sayed[1], John Mamish[2], Gonzalo Ferrer[2], Lijun Teng[1], Ross M. McKenzie[1], Alfred O. Hero[2], Edwin Olson[2] and Adam A. Stokes[1]*

[1] School of Engineering, Institute for Integrated Micro and Nano Systems, The University of Edinburgh, Edinburgh, United Kingdom, [2] Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, United States

Scaling up robot swarms to collectives of hundreds or even thousands without sacrificing sensing, processing, and locomotion capabilities is a challenging problem. Low-cost robots are potentially scalable, but the majority of existing systems have limited capabilities, and these limitations substantially constrain the type of experiments that could be performed by robotics researchers. As an alternative to increasing the quantity of robots by reducing their functionality, we have developed a new technology that delivers increased functionality at low-cost. In this study, we present a comprehensive literature review on the most commonly used locomotion strategies of swarm robotic systems. We introduce a new type of low-friction locomotion—active low-friction locomotion—and we show its first implementation in the HoverBot system. The HoverBot system consists of an air levitation and magnet table, and a HoverBot agent. HoverBot agents are levitating circuit boards that we have equipped with an array of planar coils and a Hall-effect sensor. The HoverBot agent uses its coils to pull itself toward magnetic anchors that are embedded into a levitation table. These robots use active low-friction locomotion; consist of only surface-mount components; circumvent actuator calibration; are capable of odometry by using a single Hall-effect sensor; and perform precise movement. We conducted three hours of experimental evaluation of the HoverBot system in which we observed the system performing more than 10,000 steps. We also demonstrate formation movement, random collision, and straight collisions with two robots. This study demonstrates that active low-friction locomotion is an alternative to wheeled and slip-stick locomotion in the field of swarm robotics.

Keywords: HoverBot, swarm robots, design for manufacturability, low-friction locomotion, precise locomotion, robot testbed, physical simulation

INTRODUCTION

Swarm robotics is the study of developing and controlling scalable groups of simple robots. Individual robots within a swarm only possess limited capabilities. They move in two- or three-dimensional space, sense their local environment, and communicate with only their nearest neighbors. These local interactions between hundreds or thousands of robots can potentially give rise to complex behaviors (Brambilla et al., 2013). Much swarm robotics research is inspired by the observation of emergent behaviors in nature (Bonabeau et al., 1999). Colonies of termites work together to build complex structures that are of great importance for survival of the colony as a whole. Schools of fish cluster together making it difficult for a visually orientated predator to pick and grab an individual before

it disappears into the school. Flocks of birds fly in formation to utilize the flapping of the front bird's wing, which creates uplift and eases locomotion for the remaining flock. Control in these three natural systems is entirely distributed among the individuals, without having a leader coordinating activities. These natural systems accomplish complex global tasks through simple local interactions of large groups of autonomous individuals and are commonly referred to as examples of swarm intelligence.

Much research in swarm robotics has been conducted *via* computer simulations. Brambilla et al. analyzed more than 60 publications that dealt with swarm robotic collective behaviors in 2013. They found that more than half of these publications presented results which were obtained through simulations or models (Brambilla et al., 2013). Although simulators are a valuable tool for systematically exploring the algorithmic-behavior of a swarm, they frequently involve simplifications and reductionist axioms to enable computational tractability. Such simulated systems can fail to faithfully reproduce the intricate physical interactions and variability that exist in real systems, and their fidelity to the real world is difficult to verify or improve without feedback from physical experiments (Rubenstein et al., 2014).

Building physical systems, however, is a challenging task. Swarm robotics researchers frequently face a cost-functionality optimization problem when it comes to building a scalable robot swarm. For example, every additional sensor on a robot increases the power consumption of the system, requires an additional sensor specific input on the microcontroller, requires additional space, and increases the overall cost. As a result, research in large-scale swarms (>1,000) often sacrifices sensing, processing and locomotion capabilities for the size and quantity of robots, and these design decisions substantially limit the type of experiments that researchers can perform. Instead of increasing the quantity of robots in a swarm by reducing the functionality of each robot, the robotics community requires new technologies that deliver increased functionality at low cost.

Motivation

Our work on technologies for swarm robotics is motivated by three primary objectives, we want to: decrease the cost of fabrication, ease the process of fabrication, and increase the precision of locomotion. We believe that these three factors, among many others, play a crucial role in the development of the next generation of swarm robotic systems. In addition to the obvious focus on decreasing cost, we observed that there is a considerable manufacturing-assembly overhead for existing swarm systems that use either wheeled or slip-stick locomotion. Every component on a robotic system that has to be manually assembled by the researcher invokes a labor cost. This requirement for manual labor by skilled-engineers limits the practicality of fabricating and experimenting with robot collectives at scale.

Improving movement precision enhances localization, whereas precise localization is a *useful* technology to achieve coordination and control of swarm robots (Wu et al., 2014). It is not an easy task for simple robots to maneuver precisely and to reach a common goal. Generally, the difficulties are due to hardware constraints such as small sensor ranges, very limited computational power, little memory, and imprecise locomotion (Moeslinger et al., 2011).

For example, low-cost locomotion strategies such as slip-stick locomotion suffer from imprecise movement. Vibration motors provide noisy locomotion without positional feedback, thus preventing a single robot from traveling long distances with any known precision (Rubenstein et al., 2014).

We have developed a locomotion strategy—active low-friction locomotion—that allows agents to maneuver precisely on a discrete two-dimensional grid. Its first embodiment—the HoverBot system—is easy to fabricate and to further-customize. The entire robot consists of a single printed circuit board (PCB), surface-mount components, and a battery. HoverBots can be ordered in large-number from a circuit-board manufacturer in panel-format and arrive fully populated with components—ready to use—thereby lowering the barrier to entry for researchers wishing to study complex systems using swarm robots.

Locomotion Strategies of Swarm Robotic Systems

This study briefly reviews locomotion strategies used by previous swarm robotic systems, it introduces our new locomotion strategy, and puts it into perspective against the literature. Specifically:

(1) We analyze the locomotion methods of 16 swarm robotic systems found in the literature and provide a summary in **Table 1**. The content of **Table 1** is based on the cited work shown in the first column of each row.
(2) We associate each locomotion strategy (wheeled and slip-stick locomotion) with a representative system from **Table 1**. We compare, in detail, the advantages and disadvantages of locomotion strategies in **Table 2**.
(3) We explain and demonstrate our active low-friction locomotion strategy, and we present its first implementation, the HoverBot system.

Tables 1 and **2** contain specific terminology. While most terminology for these features is self-evident, we provide here a summary for those that may be unclear. "Hardware odometry" is defined as the use of sensors to estimate change in position over time. This term indicates systems which do not possess a real form of odometry or which address the lack of hardware odometry by performing collective algorithms such as in Rubenstein et al. (2012). In this column, N/A refers to the fact that the cited publication does not explicitly state information about odometry. "Type of motion" clarifies whether a motion is continuous or discrete and if discrete with what step size. "Dependencies" refer to specific environments which the robots require to function properly. "Surface-mount-technology (SMD)" components are components which can be soldered directly onto a PCB. "Non-SMD" components are usually incompatible with pick-and-place machines and often require manual assembly which generally increases the labor effort and cost for mass manufacture.

Previous Swarm Robotic Systems

The swarm robotic systems listed in **Table 1** use either wheeled or slip-stick locomotion. Slip-stick locomotion refers to the alternation between slipping and sticking of an agent to a substrate that results into directed locomotion (Vartholomeos and

TABLE 1 | Comparison of 16 swarm robotic systems found in the literature.

Reference	Robot name	Locomotion strategy	Propulsion system	Hardware odometry
Mondada et al. (1994)	Khepera	Wheeled	DC motors	Wheel encoders
Siegwart and Caprari (2003)	Alice	Wheeled	Bidirect. motors	Wheel encoders
Mondada et al. (2003)	S-bot	Treeled	DC motors	Wheel encoders
Kornienko et al. (2005)	Jasmine	Treeled	DC motors	N/A
Seyfried et al. (2005)	I-swarm	Slip-stick	Piezoelectric Polymer	No
Mondada et al. (2006)	e-puck	Wheeled	Stepper motors	Wheel encoders
Turgut et al. (2007)	Kobot	Wheeled	DC motors	N/A
Bonani et al. (2010)	MarXbot	Treeled	Rotational motors	Accelerometer, gyroscope
Rubenstein et al. (2012)	**Kilobot**	**Slip-stick**	**Vibration motors**	**No**
McLurkin et al. (2012)	R-one	Wheeled	DC motors	Wheel encoders
Riedo et al. (2013)	Thymio II	Wheeled	DC motors	N/A
Farrow et al. (2014)	Droplet	Slip-stick	Vibration motors	No
GCtronic (2017)	Elisa-3	Wheeled	DC motors	Wheel encoders
Pickem et al. (2015)	**GRITSBot**	**Wheeled**	**Stepper motor**	**Stepper motor**
Wilson et al. (2016)	Pheeno	Wheeled	DC motors	Wheel encoders
Nemitz et al. (2017) (this paper)	**HoverBot**	**Low-friction**	**Planar coils**	**Hall-effect sensor**

The three highlighted rows depict the swarm robotic systems whose locomotion strategy is further analyzed in **Table 2**.

TABLE 2 | Comparison of wheeled, slip-stick, and low-friction locomotion.

Feature	Wheeled (Pickem et al., 2015)	Slip-stick (Rubenstein et al., 2012)	Active low-friction
Robot velocity (cm/s)	25	1	0.64
Type of motion	Continuous	Continuous	Discrete—equidistant ~20 mm steps[a]
Battery lifetime	30 min–5 h at 150 mAh	3–24 h at 160 mAh	25 min–600 h at 300 mAh
Dependencies	No	Flat surface	Levitation table with embedded magnets
Hardware odometry	Stepper motors	No	Hall-effect sensor
Actuator calibration	Not required	Required	Not required
Number of non-surface-mount components[b]	≥4	≥5	0
Difficulty of mechanical assembly[c]	(1)(2)(3)(4)(5)(6)	(1)(3)(6)	(6)
Cost[d] at 1,000 units ($)	13.34	3.12	1.96

[a]Robust (error tolerant) movement on a discrete grid is equivalent to precise movement.
[b]Wheeled: two wheels, two motors, and motor control board. Slip-stick: three legs, two vibration motors, and electronics. Active low-friction: electronics.
[c](1) Soldering non-surface-mount components, (2) cutting components, (3) gluing components, (4) screwing components, (5) stacking components, and (6) connecting battery.
[d]Cost for components that are solely associated with locomotion, in order quantities of 1,000.

Papadopoulos, 2006). The vast majority of swarm robotic systems use wheeled locomotion with DC motors and wheel encoders. There are a few exceptions which use tracks and wheels (treels) and accelerometers, gyroscopes, or stepper motors for odometry. Treels are considered as wheeled locomotion. Three systems use slip-stick locomotion, whereas two of those three systems use vibration motors and the remaining system uses piezoelectric polymers as actuators. The HoverBot System is the first implementation of our active low-friction locomotion.

Comparison of Locomotion Strategies

Table 2 compares wheeled, slip-stick, and low-friction locomotion by using the GRITSBot, the Kilobot, and the HoverBot as representative systems. We selected Kilobot as a representative for slip-stick locomotion because it is the first and only large-scale robot swarm exceeding a collective size of 1,000 units. We chose GRITSBot as a representative for wheeled locomotion. Pickem et al. (2015) have presented a recent system that explores both cost and functionality.

While wheeled locomotion has advantages in robot velocity, platform independence, hardware odometry, and actuator calibration, it has disadvantages in battery lifetime, number of non-SMD components (minimum two wheels and two motors), difficulty of mechanical assembly, and cost (including motor control board). In Pickem et al.'s work, non-surface-mount components had to be soldered, receiver coil wires needed to be cut and glued, wheels had to be screwed onto motors, circuit boards needed to be stacked, and the battery had to be connected.

In comparison, slip-stick locomotion has advantages in battery lifetime and cost, but disadvantages in robot velocity, the dependency on flat surfaces, hardware odometry, actuator calibration, and number of non-SMD components (the minimum number being three legs and two vibration motors). In Rubenstein et al.'s work, their mechanical assembly consisted of soldering non-surface-mount components, gluing vibration motors to the robot, and connecting a battery.

Our active low-friction locomotion has advantages in that it provides hardware odometry, requires no actuator calibration, has no non-SMD components, simple mechanical assembly, and is low cost; but it has disadvantages in robot velocity, dependency on a levitation-magnet table, and battery lifetime. To mechanically assemble our robot, one must only connect a battery.

Overall, each of the three strategies possesses specific advantages over the others.

The contribution of this study is the introduction of an active low-friction locomotion mechanism and its first embodiment, the HoverBot system. In addition to using active low-friction locomotion, the HoverBots have the following characteristics, they:

- possess odometry by using a single Hall-effect sensor;
- only require electronic components that are surface mountable;
- only require connecting a battery to a robot as an assembly step;
- use low-cost actuators and associated circuitry;
- do not require actuator calibration;
- move precisely on a discrete grid.

LOW-FRICTION LOCOMOTION

To move—on land, in water, or in the air—always requires an expenditure of energy. Reducing the resistance to motion, namely, friction, allows a greater range of travel for a given input of energy (Radhakrishnan, 1998). However, instead of enhancing locomotion, we enable locomotion by reducing friction. A good example of our proposed locomotion mechanism can be observed in nature. *Nannosquilla decemspinosa* is a small stomatopod found in sand substrates on the Pacific coast of Central and South America. These stomatopods are capable of maneuvering if supported by a 1-mm layer of water and lose this capability once their surrounding dries up (Caldwell, 1979).

The HoverBot is conceptually similar to *N. decemspinosa* and is only capable of maneuvering if it is supplied with a constant air flow beneath its contact surface. The airflow reduces the friction between robot and table allowing relatively weak forces to be used for locomotion. Specifically, we embedded permanent magnets into a levitation table. The HoverBot possesses planar coils which interact with these permanent magnets, resulting in two-dimensional locomotion. Such forces would be insufficient if friction had not been reduced. This concept relaxes actuator boundaries allowing a significant simplification of the robot's actuation and control system.

We define *active* low-friction locomotion as a locomotion type that enables robots to maneuver autonomously, and we define *passive* low-friction locomotion as locomotion type that allows robots to maneuver heteronomously.

Not included in **Table 1**, but relevant to our technical approach, is work from Groß et al. (2011), Napp et al. (2011), Cappelleri et al. (2014), and Pelrine et al. (2017). Groß et al. reported on an experimental setup in which they investigated aided assembly with floating building blocks using an air table. Their system used *passive* low-friction locomotion in which their building blocks did not possess locomotion capabilities, but modules would flow passively in the agitated medium. Napp et al. investigated stochastic interactions between active and passive robots using *passive* low-friction locomotion. Passive robots were foam blocks with complementary shape and embedded magnets that assembled over time on an air bed. Active robots, while not capable of autonomous movement, could expend energy to disassemble the passive robots. Cappelleri et al. introduced a novel approach to achieving independent control of multiple robot magnets. In their work, they designed a grid of planar microcoils. The coils were used to generate magnetic potentials to control the trajectories of magnets. Pelrine et al.'s work is similar to Cappelleri's, but differs in that they add onto their PCBs a thin graphite layer that makes their magnet robots levitate. Both their work feed into additive micromanufacturing with swarms. Similarly to Groß's and Napp's work, agents did not possess locomotive autonomy but were moved by external stimuli; all four approaches are relevant but distinctly different to the work we present here.

The Levitation–Magnet Table

Figure 1 illustrates the concept, and our implementation, of the levitation–magnet table. The table supplies an airflow beneath the HoverBots' contact surface creating an air cushion that reduces friction between the robot and the locomotion substrate. The differential pressure required to lift a HoverBot can be estimated according to Leal (2007) by the following equation:

$$\Delta P = (P_2)_{\min} - (P_{\text{amb}}) \geq \frac{M * g}{\pi * R^2}. \tag{1}$$

Equation 1 implies that an increase in the robot's weight or a reduction of the robot's surface area can be encountered by an increase in differential pressure. In our experiments, we required approximately 22.5 mm H_2O differential pressure to levitate HoverBots. We measured the differential air pressure between air chamber and ambient environment by using a u-tube manometer. We controlled the air blower's supply voltage with an adjustable transformer (Variac) which varied the air blower's output air-flow-rate, and which in-turn varied the differential pressure between the inside and outside of the levitation table.

The levitation–magnet table measures 200 mm × 300 mm and has an array of permanent magnets embedded into its surface. The permanent magnets serve a double purpose, they: (1) act as magnetic anchors that a HoverBot utilizes to maneuver and (2) give rise to a magnetic field with a discrete regular pattern of features which a HoverBot with a Hall-effect sensor can utilize for odometry. All magnets were assembled mono-directionally: north-pole facing up.

The HoverBot

A HoverBot consists of a single four-layer PCB, shown in **Figure 2**, and a detachable 300 mAh lithium polymer battery. The bottom layer comprises five planar actuation coils. Each HoverBot has a diameter of 39 mm and weighs 19.4 g with, and 7.4 g without, a battery. HoverBot possesses a low-power microcontroller, programming and debug ports, an infrared transceiver, a Hall-effect sensor, and a transistor circuit.

Actuation

We embedded the planar coils in a cross-formation into the bottom layer of the PCB. Each actuation coil has 17 turns and a trace width of 150 μm. A trace width of 150 μm and one $\frac{oz}{ft^2}$ trace thickness allows maximum currents of approximately
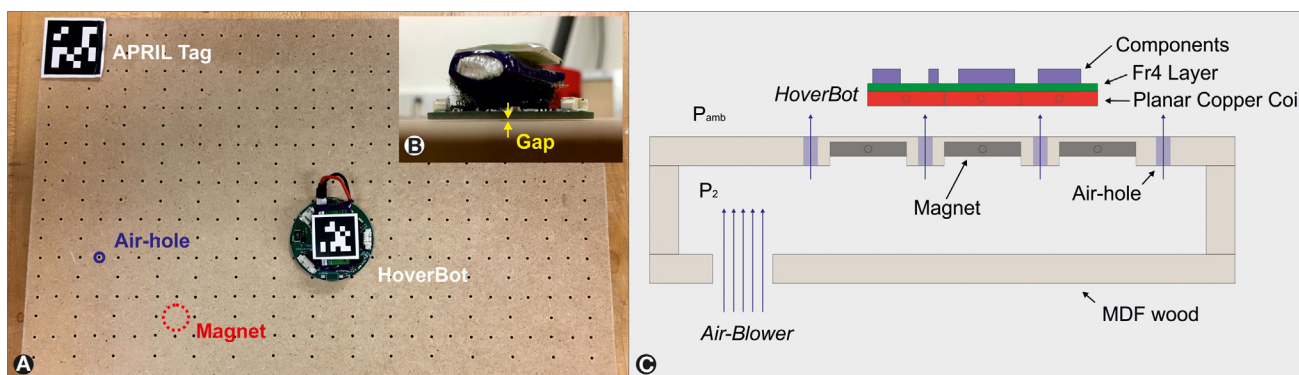
FIGURE 1 | (A) Levitation–magnet table with a HoverBot on top. There are four AprilTags for each of the table corners (one visible in the figure) and one AprilTag attached to the HoverBot. The AprilTags are used for tracking (Olson, 2011). **(B)** A photograph showing an air gap between a HoverBot and the table. **(C)** Conceptual overview: an air blower increases pressure P2 within the air chamber. The pressure difference between Pamb and P2 causes a HoverBot to levitate, hence the friction between robot and table decreases.
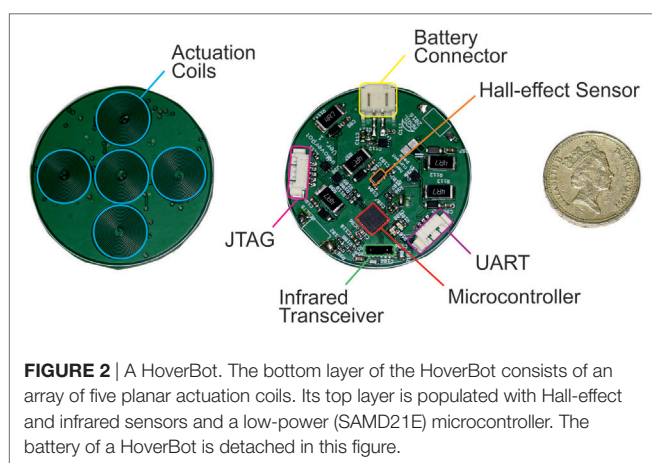
FIGURE 2 | A HoverBot. The bottom layer of the HoverBot consists of an array of five planar actuation coils. Its top layer is populated with Hall-effect and infrared sensors and a low-power (SAMD21E) microcontroller. The battery of a HoverBot is detached in this figure.

300 mA based on the Generic Standard on the Printed Board Design (IPC-2221) charts. We set the maximum current per coil to 500 mA, which induces a magnetic field of 1.1 mT. Our design uses a maximum current that is greater than the suggested standard, because we decided to evaluate the circuitry to its upper limits. We measured the magnetic field by using an InvenSense MPU-9250 magnetometer. We placed the magnetometer onto the core of the center coil.

Each coil is connected in series with a current limiting resistor and a transistor. If the transistor switches on, a constant voltage is applied across the coil and resistor. The transistor's switching behavior is controlled by a pulse-width-modulated (PWM) signal from the microcontroller. We control the amount of current through the coil by changing the duty cycle of the PWM. The magnetic field of a solenoid can be approximated by Ampere's law:

$$B = \mu * n * I, \tag{2}$$

$$n = \frac{N}{L}, \tag{3}$$

where B is magnetic flux density, μ is permeability; n is turn density; I is current; N is the number of turns; and L is unit length.

Fundamentally similar to Eqs 2–3, the magnetic field of a planar coil is dependent on the coil's turn density and the current flow. The number of turns is identical for every HoverBot. However, the coil, trace, and current limiting resistor resistance could vary due to manufacturing tolerances and cause a change of current flow for a given duty cycle. We measured the average series resistance of 15 actuation circuits of a total of 3 HoverBots with a Fluke 115 multimeter. The SD was 0.1 Ohm, which causes a current change of 7 mA. Hence, the potential current fluctuations are less than 1.5% and can be neglected. HoverBots do not require any kind of actuator calibration.

Sensing and Communication

A HoverBot possesses infrared and Hall-effect sensors. The Hall-effect sensor can be used for odometry and the detection of local magnetic fields. The infrared transceiver can be used for robot-to-computer communication. Our current HoverBot version does not allow robot-to-robot communication due to limitations in its hardware configuration. It only possesses a single IR transceiver pointing upwards.

Programming and Debugging

A HoverBot has programming and debug ports (IR transceiver, JTAG and UART). We programmed the HoverBot *via* JTAG using an Atmel SAM-ICE programmer. Therefore, this HoverBot version requires a wired connection to be programmed. We debugged HoverBot *via* infrared using an infrared handheld device.

Power System

A HoverBot has 3.7 V 300 mAh lithium polymer batteries attached to it. We calculated the minimum battery life by accumulating the currents that occur during locomotion.

The current locomotion strategy requires a constant current of approximately 720 mA, which allows a minimum battery life of around 25 min. However, lithium polymer batteries should never be completely discharged due to their chemistry. We wrote a battery-watch program to monitor the battery during runtime, and this program shuts down all circuitry when the battery reaches 90% depletion. The maximum battery life is calculated by considering HoverBot when in sleep mode, in which it approximately consumes 500 μA. In this low-power mode, the HoverBot's battery life time rises to around 600 h or 25 days. In this HoverBot version, the lithium polymer batteries have to be detached for charging. We charged the batteries by using a Turnigy Micro-6 LiPoly battery charger.

Locomotion Strategy

The HoverBot levitates on air cushions and maneuvers by sequentially energizing its planar coils to pull itself toward magnetic anchors. **Figure 3** indicates a HoverBot's open-loop locomotion strategy. A single step, a movement from one magnetic anchor to another, is decomposed into three part steps. In step 1, HoverBot starts from its idle state in which its center coil is aligned with a magnet, and the other four coils are each overlapping with adjacent magnets. The HoverBot simultaneously actuates one side coil with maximum current and the opposite side coil with medium current. This actuation results in an overall movement to the right while preventing HoverBot from rotating. Subsequent steps are conceptually the same, but each step requires a differing pair of coils to be actuated. Three of these steps are required for a HoverBot to move from one magnet to another. This actuation scheme only enables complete magnet-to-magnet movements. A change of direction during a part step has not been investigated. The relative positions of the HoverBot coils and the magnets are crucial for this actuation scheme. We chose magnet-to-magnet and coil-to-coil pitches based on Eq. 4 to ensure a 50% overlap between actuator coil and an adjacent magnet at any given step assuming that coil and magnet diameters match. Therefore,

HoverBot's minimum step size is the pitch between adjacent magnets (2 cm pitch).

$$r_{mc,c} = \frac{d_{mc}}{d_c} = \frac{d_m - d_c}{d_c} = \frac{1}{2}, \tag{4}$$

where d_m is magnet to magnet pitch; d_c is coil to coil pitch; d_{mc} is magnet to coil pitch; and $r_{mc,c}$ is ratio of d_{mc} to d_c.

A HoverBot moves in a two-dimensional discrete environment. The programmer cannot deliberately rotate a HoverBot or move it in any other trajectories than the Manhattan Geometry.

Odometry

Figure 4 is based on the Hall-effect sensor readings from a HoverBot during movement which were paired with spatial information from the AprilTags. While a HoverBot moves from one magnetic anchor to another (2 cm pitch), its Hall-effect sensor measures a continuously changing magnetic flux density as indicated in **Figure 4**. The Hall-effect sensor is centered above the center coil and is capable of measuring magnetic flux densities from −73 to +73 mT. The maximum readings occur when the HoverBot's center coil is aligned with a magnetic anchor. Although our current actuation scheme operates as an open-loop control, the magnetic flux density changes over distance depict distinct features in the two-dimensional space which could be utilized as feedback for closed-loop control. We have not experienced distorted sensor (Hall-effect, IR) readings due to magnetic interference.

Manufacture and Cost

The circuitry of a HoverBot only consists of surface-mount components as indicated in **Table 2**. HoverBot is designed explicitly for manufacturability; it consists of a single PCB and therefore mass manufacture is a simple case of placing a batch order with a PCB foundry. HoverBots can be autonomously populated with pick-and-place machines at the point of manufacture. Assembly of one robot takes seconds since it only consists of plugging in a

FIGURE 3 | The illustrated actuation scheme describes a movement to the right. The other directions can be derived from this actuation profile. A HoverBot moves from one magnetic anchor to another by performing three steps. Starting in idle state (step 1), a HoverBot always actuates two coils at the time for any given step since we require two points in space to maintain the orientation of a plain. Therefore, this actuation sequence prevents rotation during movement.

FIGURE 4 | Experimental results showing that the magnetic flux density peaks correspond to the centers of the permanent magnets. The magnets are 2 cm displaced from one another. A HoverBot in idle state aligns with its Hall-effect sensor above a magnet. A step from one magnet to another causes the Hall-effect sensor to measure a peak-to-peak magnetic field curve. Specifically, a movement is decomposed into three individual steps in which the Hall-effect sensor measures a distinct magnetic flux density after each step, as indicated by a black dot.

TABLE 3 | Cost summary for HoverBots in order quantities of 15 units, and for one levitation–magnet table.

Levitation–magnet table		HoverBots at 15 U	
Category	Cost ($)	Category	Cost ($)
Wood	20	Locomotion	4.10
Air Blower	55	Computation	2.05
Variac	135	Power	4.96
Magnets (140)	25	Sensing	1.09
		Communication	4.28
		Miscellaneous	3.23
		Printed circuit board manufacture	2.66
Total	235	Total	22.37

battery to a HoverBot, also indicated in **Table 2**. The fabrication of the levitation–magnet table is described in detail in Section "Fabrication."

Table 3 summarizes the costs of the current levitation–magnet table and HoverBots in order quantities of 15. The levitation–magnet table costs $235 whereas each HoverBot costs $22.37 in quantities of 15 and $11.88 in quantities of 1,000 s. The most expensive part of the levitation–magnet table is the variable transformer at $135. The costs for components that are solely associated with HoverBot's actuation system (transistors, shunt resistors, diodes, and capacitors) are $1.96 in order quantities of 1,000 as indicated in **Table 2**.

The bill of materials, HoverBot system design files, and code are available on request.

EVALUATION

To evaluate the HoverBot system we designed a controllable experimental setup. We used artificial features (fiducials)—AprilTags (Olson, 2011)—which we placed on top of the HoverBot and at each corner of the table. AprilTags are robust to occlusions and lens distortion while being very efficient in achieving detection rates of 20 Hz in our setting. To measure the accuracy of the HoverBot, we tracked the centroid and the orientation of the robot by detecting the corresponding AprilTag. **Figure 5** depicts the main features of the tracking system. This system can run for hundreds of minutes without human intervention, thereby automating the data acquisition pipeline. We used a Chameleon 1.3 MP Color (Sony IXC445) camera and a Tamron 13FM28IR 2.8 mm f/1.2 day/night lens.

We tested the HoverBot system and its low-friction locomotion by conducting eleven experiments that lasted a total of 3 h and more than 10,000 steps. In these experiments, the HoverBot circled on an arbitrary trajectory until it was nearly discharged. We used a set of AprilTags to track the HoverBot over time and subsequently evaluated its distance traveled, velocity, and number of missteps (errors). With our current actuation sequence, the HoverBot moves an average of 0.64 cm/s with an SD of

0.015 cm/s. We did not observe any missteps or accidental rotations during these three hours. A "misstep" is defined as an unsuccessful series of energized coils that results in the robot staying on its previous position. An "accidental rotation" is defined as an inadvertent robot rotation by 45° due to local table imperfections (e.g., air flow fluctuations) or collisions with other robots or static objects. A video recording of this experiment is provided by the Supplemental Video—SV1.[1]

Although the HoverBot moves robustly, we observed unintentional shaking in all four directions during movement. Therefore, we compared moved distance, which includes the total distance

FIGURE 5 | Experimental setup to evaluate a HoverBot's locomotion performance. We placed one AprilTag in each corner of the levitation–magnet table. These tags serve as reference points and allow determination of a HoverBot's relative position over time. Each AprilTag corresponds to an ID number. During experiments, we read out each AprilTag's ID, x-position, y-position, and rotation. The red line indicates a HoverBot's trajectory, which reinforces with each lap.

including unintentional shaking, with effective distance, which is the actual distance between waypoints. We define ε by the following term:

$$\epsilon = \frac{\text{moved distance}}{\text{effective distance}}. \qquad (5)$$

We found that ε is 2.29 (on average) with an SD of 0.27. An ε of 2.29 explicitly states that the HoverBot moves more than two times the distance it travels. ε is directly related to the actuation scheme. We chose an actuation scheme that is relatively slow, but very robust by performing zero missteps over three hours of experiments. ε can be further reduced by changing the actuation scheme and specifically the timing and amount of current that flows through up to five coils simultaneously.

Recovery from a Locked Rotational Position

Although we have not experienced any accidental rotation incidents during three hours of testing, we developed an actuation strategy that allows a HoverBot to recover from a locked position. As shown by **Figure 4**, the Hall-effect sensor measures a local magnetic minimum if a HoverBot is locked due to accidental rotation. When a HoverBot recognizes this state, it can execute a recovery actuation scheme. It first actuates only its center coil to change from the position in **Figure 6A** to the position in **Figure 6B**. Then it actuates, in addition, a side coil to regain the correct orientation as indicated in **Figure 6C**. We recorded this sequence, a video–recording of this experiment is provided by the Supplemental Video—SV2.[2]

DEMONSTRATION

In addition to our quantitative evaluation of HoverBot's locomotion capabilities, we performed four additional demonstrations to give more insights into the nature of the HoverBot
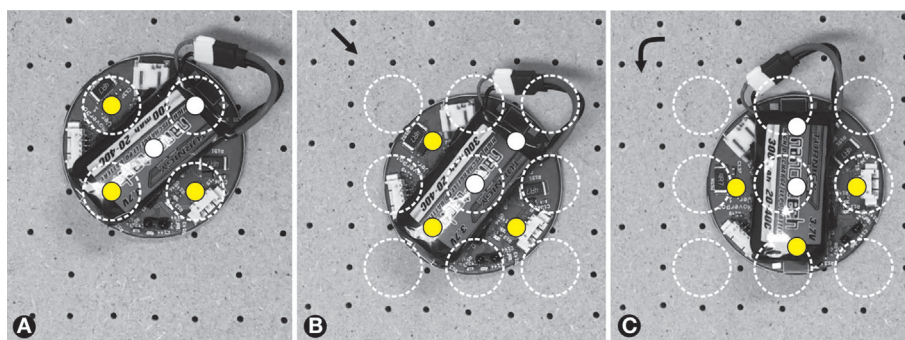
FIGURE 6 | Recovery in locked position. **(A)** A HoverBot is locked in a 45°-angled position. Four of its five coils are aligned with permanent magnets. **(B)** The HoverBot rapidly pulsed (only) its center coil and regained center coil alignment with a permanent magnet. **(C)** The HoverBot additionally actuated a side coil and regained a slightly shifted idle position. At this stage, the HoverBot is reenabled to move.

system. **Figure 7** indicates two HoverBots moving in formation (A), moving randomly, colliding, and recovering from rotation (B), colliding (C), and colliding while one robot is in sleep mode, acting as a passive agent (D).

We observed that two HoverBots that move independently in formation become unsynchronized over time due to oscillator imperfections. Physical inter-robot interactions can either result in robots maintaining their position after collision, which likely happens in a frontal collision event, or *robots* loose orientation and have to recover. The random collision demonstration also indicates possible orientation loss due to rapid and constant change in direction. Those incidents, however, are scarce, they were detected, and were recovered from. In most cases, moving HoverBots are capable of pushing passive agents, sliding them to one side, or pushing them in front, in the direction of travel. However, we recorded one incident in which a moving agent could not pass a passive agent due to a specific physical orientation. Video recordings of the experiments which correspond to **Figures 7A–D** are provided by the Supplemental Videos—SV3–SV6.[3]

DISCUSSION

Battery Life and Robot Velocity

HoverBot possesses a relatively short battery lifetime (~25 min) due to high coil actuation currents that are required to achieve magnetic fields of approximately 1.1 mT. According to Eqs 2 and 3, the magnetic field is linearly dependent on the actuation current, but also on the number of coil windings. An increase of coil windings as well as the stacking of planar coils (multilayer PCBs) could significantly decrease the power consumption.

The existing robot velocity can be improved without an increase of power consumption. The product of current and time for slow coil actuation does not change for rapid coil actuation. HoverBot's velocity is currently slow because we wanted to start off with a robust actuation scheme. Future work will have to investigate faster actuation schemes. It is very likely that actuator calibration will become necessary once we reach HoverBot's physical speed limits. The actuation schemes will become more delicate and have to energize the actuation coils extremely precisely, both in terms of the amount, duration, and direction of current flows. One solution to this control problem could be to use machine learning algorithms. An external camera system could send feedback to the HoverBot agent and inform the controller whether movement was successful or not.

Ease of Robot Fabrication

Although HoverBots only consist of surface-mount components, we believe the importance of this advantage will decrease over time. The current state of swarm robotics research requires low-cost, easy-to-fabricate, and easy-to-use swarm robotic systems. However, once we obtain a better understanding of complex systems and how emergence occurs, cost and ease of fabrication will become secondary because the risk-factors involved in deploying swarms (system failure, loss of control, and safe and reliable operation) will have decreased. Furthermore, there are many great examples in industry in which very sophisticated products have been mass manufactured (computers, cars, airplanes, etc.). Investing into an expensive swarm of robots will become worthwhile once we know how to safely operate and control it.

The Table

The existing ratio of magnet-to-magnet and coil-to-coil distances was chosen to simplify HoverBot's actuation circuitry by only requiring coils to be energized in one direction. In future work, we can investigate the use of H-bridge drivers to improve locomotion by allowing bidirectional currents to energize the actuator coils. There may also be a benefit of designing different magnet patterns, such as those which that vary between polarities as well as exploit different geometric constellations (e.g., concentric patterns).

Scaling the System

The current table measures 200 mm × 300 mm, and this size limits the maximum number of robots on the table to 35, assuming a lattice robot formation without extra space for movement and a robot diameter of 40 mm. There is no reason why the table or robot could not be scaled in either direction. The table size could be significantly increased, to the size of an air hockey table for example. The differential pressure that causes the robots to levitate can be easily increased by using a more powerful blower, or even several at once. An increase in differential pressure would allow greater payloads to be carried by the robots. The robot size

FIGURE 7 | Demonstrations of the locomotion capabilities of multiple HoverBots. **(A)** Two HoverBots circle in formation until they are unsynchronized—video SV1. **(B)** Two HoverBots move randomly, collide, and recover—video SV2. **(C)** Two HoverBots collide frontally with one another—video SV3. **(D)** One HoverBot collides with a passive HoverBot—video SV4. Red and blue trajectories depict the HoverBot's movements over time.

could be significantly increased or decreased. There are micromachining systems that are capable of fabricating 50 µm wide copper traces (e.g., LPKF Protolaser U3) allowing much smal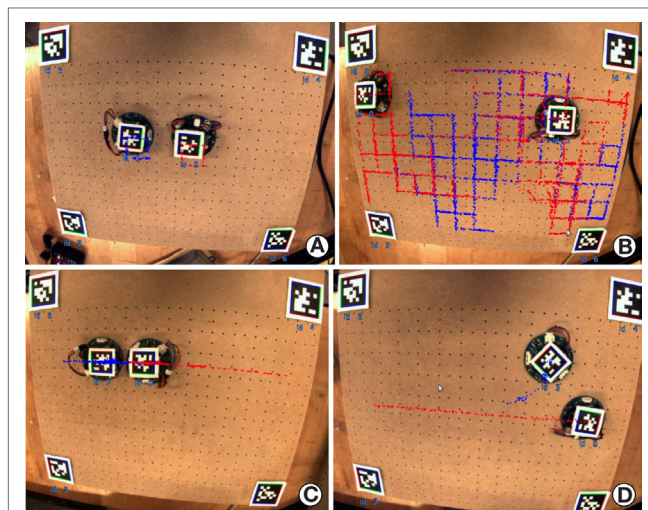ler actuator coil sizes. The 300 mAh battery could be substituted with less powerful batteries or even replaced with solar panels.

Future Directions

HoverBot version-2 should possess four directional communications to increase further its utility as swarm algorithmic testbed.

The collision of an active with a passive robot in video SV6 indicate an opportunity for new swarm robotic algorithms in which passive and active robots are being utilized to achieve a task. A passive robot might become active if it has not been pushed around by another robot for a defined period of time. A passive robot might also specialize in sensing and inform active robots about its observations. This heterogeneity might lead to strategies that optimize the power budget of the swarm while solving the task at hand.

The formation demonstration in video SV3 indicates that the HoverBot system can be used for even larger collective movements. This behavior is difficult to achieve with wheeled or slip-stick actuated swarm systems since such systems move in continuous space and must rotate to change directions. HoverBots locomotion can be compared with that of quadrotors in formation flight (Kushleyev et al., 2013), maintaining orientation of the local and global directions.

Almost all of HoverBot's advantages originate from its minimalist design. HoverBots levitate, move precisely on a discrete grid, and are capable of verifying a step by measuring continuously magnetic flux densities. We will utilize this combination of discrete motion with continuous local perception to study search and tracking as well as mapping algorithms. An excellent starting point is Senanayake review on search and tracking algorithms for swarm robots (Senanayake et al., 2014).

CONCLUSION

In this study, we introduced a new locomotion strategy, active low-friction locomotion, and showed its first embodiment: the HoverBot system. We demonstrated HoverBot's capabilities by performing six different experiments ranging from moving in a predetermined trajectory, to random movement and inter-robot collisions. Active low-friction locomotion is an alternative to wheeled- and slip-stick locomotion in the field of swarm robotics. The HoverBot system possesses odometry by using a single Hall-effect sensor, it only requires components that are surface mountable, it only requires connecting a battery as assembly step, it uses low-cost actuators and associated circuitry, does not require actuator calibration, and moves precisely on a discrete grid. The HoverBot systems offer a unique combination of discrete precise motion with continuous local perception. Its hardware can be easily extended with additional sensors. Potential research directions using this embodied-simulation system will include search and tracking, or mapping with robot swarms.

The HoverBot system serves as a testbed for new hardware and algorithms.

FABRICATION

Fabrication of Levitation–Magnet Table

We purchased 10 mm wide and 3 mm thick cylindrical N42 magnets from Amazon. We bought 12.7 mm thick medium-density fiberboard from a local hardware store. We used a ShopBot Buddy to mill and drill holes. We used a 0.063″ drill bit for the air-holes and a 0.394″ end-mill for the magnet pockets. We placed the top-plate of the air table on an optics (metal) table and embedded the magnets mono-directionally (polarity) into the pockets. We used an Arrow TR400 glue gun to fix the magnets in the pockets. We used a Mcculloch MCB2205 leaf blower as the air source in combination with a Circuit Specialists 16VA520T20 Variac for airflow control. The air blower listed in **Table 3** is the Black & Decker BV5600 High Performance Blower (for price reference) and is equivalent to the MCB2205. We leveled the levitation–magnet table using a water scale.

Fabrication of a HoverBot

We purchased all electronics components from Digikey. The circuit boards were designed with CadSoft Eagle and manufactured by 4PCB.com. We soldered the components by using a hot air pencil and an airbath preheating system from Zephyrtronics.

AUTHOR CONTRIBUTIONS

MN: created the system and is lead author of all sections of the work. MS: contributed building the HoverBot table, developing the experiment scheme, and revised the manuscript. JM: contributed building the HoverBot agent. GF: contributed to the AprilTag setup, data analysis, revising and partly writing the manuscript. LT and RM: contributed to the manuscript/revision and the development of the table. EO and AH: advised on building the HoverBot system. AS: lead advisor and primary editor of the manuscript.

ACKNOWLEDGMENTS

FUNDING

REFERENCES

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*, Vol. 1. New York, NY: Oxford University Press.

Bonani, M., Longchamp, V., Magnenat, S., Rétornaz, P., Burnier, D., Roulet, G., et al. (2010). "The marXbot, a miniature mobile robot opening new perspectives for the collective-robotic research," in *IEEE/RSJ 2010 International Conference on Intelligent Robots and Systems, IROS 2010 – Conference Proceedings*, Taipei, 4187–4193.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.* 7, 1–41. doi:10.1007/s11721-012-0075-2

Caldwell, R. L. (1979). A unique form of locomotion in a stomatopod – backward somersaulting. *Nature* 282, 71–73. doi:10.1038/282071a0

Cappelleri, D., Efthymiou, D., Goswami, A., Vitoroulis, N., and Zavlanos, M. (2014). Towards mobile microrobot swarms for additive micromanufacturing. *Int. J. Adv. Robot. Syst.* 11(9), 150. doi:10.5772/58985

Caprari, G., and Siegwart, R. (2003). "Design and control of the mobile micro robot alice," in *Proceedings of the 2nd International Symposium on Autonomous Minirobots for Research and Edutainment AMiRE 2003* (Brisbane: CITI), 22–32.

Farrow, N., Klingner, J., Reishus, D., and Correll, N. (2014). "Miniature six-channel range and bearing system: algorithm, analysis and experimental validation," in *Proceedings – IEEE International Conference on Robotics and Automation*, Hong Kong, 6180–6185.

GCtronic. (2017). *Elisa-3*. Available at: http://www.gctronic.com/doc/index.php/Elisa-3

Groß, R., Magnenat, S., Küchler, L., Massaras, V., Bonani, M., and Mondada, F. (2011). "Towards an autonomous evolution of non-biological physical organisms," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Budapest, Vol. 5777 LNAI, 173–180.

Kornienko, S., Kornienko, O., and Levi, P. (2005). "Minimalistic approach towards communication and perception in microrobotic swarms," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, Edmonton, 4005–4011.

Kushleyev, A., Mellinger, D., Powers, C., and Kumar, V. (2013). Towards a swarm of agile micro quadrotors. *Auton. Robots* 35, 287–300. doi:10.1007/s10514-013-9349-9

Leal, L. G. (2007). *Advanced Transport Phenomena: Fluid Mechanics and Convective Transport Processes*. Cambridge University Press, 934.

McLurkin, J., Lynch, A. J., Rixner, S., Barr, T. W., Chou, A., Foster, K., et al. (2012). "A low-cost multi-robot system for research, teaching, and outreach," in *Springer Tracts in Advanced Robotics*, Vol. 83 STAR, Lausanne, 597–609.

Moeslinger, C., Schmickl, T., and Crailsheim, K. (2011). "A minimalist flocking algorithm for swarm robots," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 5778 LNAI, Budapest, 375–382.

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., et al. (2009). "The e-puck, a robot designed for education in engineering," in *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*, Vol. 1, (IPCB, Instituto Politécnico de Castelo Branco), 59–65. No. LIS-CONF-2009-004.

Mondada, F., Franzi, E., and Ienne, P. (1994). Mobile Robot Miniaturisation: A Tool for Investigation in Control Algorithms. *Exp. Robot.* 501–513.

Mondada, F., Guignard, A., and Bonani, M., Bar, D., Lauria, M., and Floreano, D. (2003). "SWARM-BOT: from concept to implementation," in *Proceedings.*

2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2003. (IROS 2003), Vol. 2. (IEEE), 1626–1631.

Napp, N., Burden, S., and Klavins, E. (2011). Setpoint regulation for stochastically interacting robots. *Auton. Robots* 30, 57–71. doi:10.1007/s10514-010-9203-2

Olson, E. (2011). "AprilTag: a robust and flexible visual fiducial system," in *Proceedings – IEEE International Conference on Robotics and Automation*, Shanghai, 3400–3407.

Pelrine, R., Hsu, A., Cowan, C., and Wong-Foy, A. (2017). "Multi-agent systems using diamagnetic micro manipulation—from floating swarms to mobile sensors," in *International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)* (Seattle, WA: IEEE), 1–6.

Pickem, D., Lee, M., and Egerstedt, M. (2015). "The GRITSBot in its natural habitat – a multi-robot testbed," in *Proceedings – IEEE International Conference on Robotics and Automation*, 4062–4067.

Radhakrishnan, V. (1998). Locomotion: dealing with friction. *Curr. Sci.* 74, 826–840. doi:10.1073/pnas.95.10.5448

Riedo, F., Chevalier, M., Magnenat, S., and Mondada, F. (2013). "Thymio II, a robot that grows wiser with children," in *Proceedings of IEEE Workshop on Advanced Robotics and Its Social Impacts, ARSO*, Tokyo, 187–193.

Rubenstein, M., Ahler, C., and Nagpal, R. (2012). "Kilobot: a low cost scalable robot system for collective behaviors," in *Proceedings – IEEE International Conference on Robotics and Automation*, St. Paul, MN, 3293–3298.

Rubenstein, M., Cornejo, A., and Nagpal, R. (2014). Programmable self-assembly in a thousand-robot swarm. *Science* 345, 795–799. doi:10.1126/science.1254295

Senanayake, M., Senthooran, I., Barca, J. C., Chung, H., Kamruzzaman, J., and Murshed, M. (2014). Search and tracking algorithms for swarms of robots: a survey. *Rob. Auton. Syst.* 75, 422–434. doi:10.1016/j.robot.2015.08.010

Seyfried, J., Szymanski, M., Bender, N., Estaña, R., Thiel, M., and Wörn, H. (2005). The I-SWARM project: intelligent small world autonomous robots for micro-manipulation. *Swarm Robot.* 3342, 70–83. doi:10.1007/978-3-540-30552-1_7

Turgut, A. E., Goekce, F., Celikkanat, H., Bayindir, L., and Sahin, E. (2007). Kobot: a mobile robot designed specifically for swarm robotics research. *Architecture* 32–42.

Vartholomeos, P., and Papadopoulos, E. (2006). "Analysis, design and control of a planar micro-robot driven by two centripetal-force actuators," in *Proceedings – IEEE International Conference on Robotics and Automation*, Orlando, FL, 649–654.

Wilson, S., Gameros, R., Sheely, M., Lin, M., Dover, K., Gevorkyan, R., et al. (2016). Pheeno, a versatile swarm robotic research and education platform. *IEEE Robot. Autom. Lett.* 1, 884–891. doi:10.1109/LRA.2016.2524987

Wu, H., Qu, S., Xu, D., and Chen, C. (2014). Precise localization and formation control of swarm robots via wireless sensor networks. *Math. Probl. Eng.* 2014:942306. doi:10.1155/2014/942306

frontiers
in Robotics and AI

Rescuing Collective Wisdom when the Average Group Opinion Is Wrong

Andres Laan[1]*, Gabriel Madirolas[1,2] and Gonzalo G. de Polavieja[1]*

[1]*Champalimaud Neuroscience Programme, Champalimaud Center for the Unknown, Lisbon, Portugal, [2]Instituto Cajal, Consejo Superior de Investigaciones Científicas, Madrid, Spain*

The total knowledge contained within a collective supersedes the knowledge of even its most intelligent member. Yet the collective knowledge will remain inaccessible to us unless we are able to find efficient knowledge aggregation methods that produce reliable decisions based on the behavior or opinions of the collective's members. It is often stated that simple averaging of a pool of opinions is a good and in many cases the optimal way to extract knowledge from a crowd. The method of averaging has been applied to analysis of decision-making in very different fields, such as forecasting, collective animal behavior, individual psychology, and machine learning. Two mathematical theorems, Condorcet's theorem and Jensen's inequality, provide a general theoretical justification for the averaging procedure. Yet the necessary conditions which guarantee the applicability of these theorems are often not met in practice. Under such circumstances, averaging can lead to suboptimal and sometimes very poor performance. Practitioners in many different fields have independently developed procedures to counteract the failures of averaging. We review such knowledge aggregation procedures and interpret the methods in the light of a statistical decision theory framework to explain when their application is justified. Our analysis indicates that in the ideal case, there should be a matching between the aggregation procedure and the nature of the knowledge distribution, correlations, and associated error costs. This leads us to explore how machine learning techniques can be used to extract near-optimal decision rules in a data-driven manner. We end with a discussion of open frontiers in the domain of knowledge aggregation and collective intelligence in general.

Keywords: collective intelligence, collective behavior, majority voting rule, machine learning, decision-making, statistical decision theory

1. INTRODUCTION

Decisions must be grounded on a good understanding of the state of the world (Green and Swets, 1988). Decision-makers build up an estimation of their current circumstances by combining currently available information with past knowledge (Kording and Wolpert, 2004; Körding and Wolpert, 2006). One source of information is the behavior or opinions of other agents (Dall et al., 2005; Marshall, 2011). Decision-makers are, thus, often faced with the question of how to best integrate information available from the crowd. Over the past 100 years, many studies have found that the average group opinion often provides a remarkably good way to aggregate collective knowledge.

Collective knowledge is particularly beneficial under uncertainty. We look to the many rather than the few when individual judgments turn out to be highly variable. Pooling opinions can then improve the reliability of estimates by cancelation of independent errors (Surowiecki, 2004;

Hong and Page, 2008; Sumpter, 2010; Watts, 2011). A seminal case study of the field concerns the ox-weighting competition reported by Galton (1907). In a county fair, visitors had the opportunity to give their guesses regarding the weight of a certain ox. After the ox had been slaughtered and weighed, Galton found that the average opinion (1198lb) almost perfectly matched the true weight of the ox (1197lb) despite the fact that individual opinions varied widely (from below 900 to above 1,500). Numerous other studies have reported similar effects for other types of sensory estimation tasks as well as other types of problems like making economic forecasts (Lorge et al., 1958; Treynor, 1987; Clemen, 1989; Krause et al., 2011).

Sometimes, rather than estimating the numeric value of a quantity, the group needs to choose the best option among a set of alternatives. In such cases, the majority vote can be seen as the analog of averaging. The majority vote can produce good decisions even when individual judgment is fallible (Hastie and Kameda, 2005). This case was mathematically analyzed in the 18th century by Marquis de Condorcet (Condorcet, 1785; Boland, 1989). Condorcet imagined a group of people voting on whether or not a particular proposition is true. Condorcet thought individuals were fallible—each individual had only a probability p of getting the answer right. Condorcet found that if p is greater than 0.5 and all individuals vote independently, then the probability that the majority in a group of N people get the answer correct is higher than p. In fact, as N grows larger, the probability of a correct group decision rapidly approaches certainty. In other words, the group outperforms the individual.

If the assumptions of Condorcet's theorem are not satisfied, then relying on the majority vote can be dramatically worse than using the opinion of a single randomly selected individual (Kuncheva et al., 2003). A similar argument can be made for relying on the crowd average in the case of making quantitative estimates. On the one hand, there are known sets of scenarios where opinion averaging clearly helps (Galton, 1907; Surowiecki, 2004; Hong and Page, 2008). While we cannot guarantee the convergence of the group average to the truth for the continuous case, we can guarantee that the distance between the truth and the average group opinion (the error) is always equal to or smaller than the average error of an individual opinion (Larrick and Soll, 2006). In this sense, the group average is guaranteed to outperform the individual.

More generally, we can measure the penalties induced by our answers in more complex ways than by simply calculating the distance between our answer and the truth. A mathematical tool known as a cost function specifies the penalties we incur for every possible combination of the truth and our answer which may occur. As previous authors have emphasized, if we measure our cost using convex mathematical functions, then, according to Jensen's inequality (Larrick et al., 2003; Kuczma and Gilányi, 2009), the crowd mean is expected to outperform a randomly selected individual. In section 5 of our review, we will provide the reader with an introduction to cost functions and Jensen's inequality and argue, as others have done (Taleb, 2013; LeCun et al., 2015), that real-world cost functions are not restricted to be convex. For non-convex cost functions, the guarantee of Jensen's inequality no longer holds, and the average group opinion can perform worse in expectation than a randomly chosen individual. Averaging methodologies, thus, sometimes lead to what might be called negative collective intelligence, where individuals outperform the collective.

When the majority vote and the average opinion fail or prove suboptimal, we can resort to other means of opinion aggregation. We will review many alternatives including the full vote procedure, opinion unbiasing, wisdom of the resistant, choosing rather than averaging, and wisdom of select crowds (Soll and Larrick, 2009; Ward et al., 2011; Mannes et al., 2014; Madirolas and de Polavieja, 2015; Whalen and Yeung, 2015), which have all been successfully used to rescue collective wisdom when more traditional methods proved unsuccessful. While the applicability of these methods is more domain dependent than the applicability of averaging strategies, practice has shown them to yield sufficiently large improvements to make their application a worthwhile endeavor. Throughout the article, we will review the more recent methodologies in the light of signal detection theory (Green and Swets, 1988) to explain when and why the newer generation of methodologies are likely to work. We will also provide new mathematical perspectives on old results such as Condorcet theorem and explain how our mathematical treatment facilitates the analysis of some simple extensions of classical results.

Recent technological advances have also opened up the possibility of gathering very large datasets from which collective wisdom can be extracted (Sun et al., 2017). Large datasets allow researchers to consider and reliably test increasingly complex methodologies of opinion aggregation. These models are often represented as machine learning rules of opinion aggregation (Dietterich, 2000; Rokach, 2010; Polikar, 2012). In the final part of our article, we review how machine learning methods can expand on more traditional heuristics to either verify the optimality of existing heuristics or propose new heuristics in a data-driven manner.

Before we proceed, it is important to note a few caveats. First, there may be reasons to use (or not use) averaging procedures which are unrelated to the problems of reducing uncertainty or the search for an objective truth. For example, Conradt and Roper (2003) have presented a theoretical treatment where the majority vote emerges as a good solution to the problem of resolving conflicts of interest within a group (such applications may in turn suffer from other problems such as the absence of collective rationality (List, 2011)). These issues remain outside the scope of the present review.

Second, many natural and artificial systems from amoebas (Reid et al., 2016) to humans (Moussaïd et al., 2010) need to implement their decision rules through local interaction rules, especially when the collectives have a decentralized structure. We will occasionally make reference to how some algorithms are implemented in distributed systems. But we are primarily interested in what can in principle be achieved by optimal information aggregators that have access to all the relevant information in the collective. Hence, considerations relating to decentralized implementations with local interactions are not our focus and also remain mostly outside of the scope of the present review. We refer the interested reader to dedicated review articles on this topic

(Bonabeau et al., 1999; Couzin and Krause, 2003; Garnier et al., 2007; Vicsek and Zafeiris, 2012; Valentini et al., 2017).

2. A BRIEF PRIMER ON STATISTICAL DECISION THEORY

We begin our review of collective intelligence with a brief survey of statistical decision theory (Green and Swets, 1988; Bishop, 2006; Trimmer et al., 2011). Statistical decision theory studies how to find good solutions to a diverse array of problems which span the gamut from everyday sensory decision-making (e.g., using both your eyes and your ears to localize the source of an external event (Stein and Stanford, 2008)) all the way to rare technocratic decision-making (e.g., using multiple risk metrics to evaluate the disaster premiums on a public building). In all these cases, one is faced with multiple useful but imperfect information sources which one has to combine in order to arrive at the final decision. It is easy to see how the aforementioned concepts relate to collective decision-making. After all, an opinion is just another information source, often useful, but sometimes fallible, and a group of opinions is merely a term used to represent the multiplicity of such information sources (Dall et al., 2005).

Statistical decision theory examines the factors that influence how to arrive at a decision in a way that makes optimal use of all the available information. In particular, it has highlighted three critical factors which need to be examined for the purposes of specifying an optimal decision rule. These relevant factors are:

1. the relation between an information source and the truth,
2. the relations that multiple information sources have between each other,
3. the cost induced by errors (deviations from truth).

We will first give an informal explanation of each factor separately and then cover applications to collective decision-making in more detail.

The relation between an information source and the truth speaks to how much information one variable carries about another variable. The mathematical characterization is usually done in terms of probability distributions and is perhaps most easily understood in the context of categorical questions. We might consider a scenario where a doctor is asked to judge 100 medical images regarding whether or not they depict a cancerous mole. Provided we have determined which images contain cancerous moles through an independent means (perhaps by using histological techniques), we can calculate the accuracy of the doctor by computing the percentage of cases where the doctor gave an opinion coinciding with the truth. This number acts as an estimate of how likely it is for the doctor to give the correct diagnosis when she is asked to evaluate a new case.

We can gain even further insight into the doctor's performance by examining the idea of confusion matrices (Green and Swets, 1988; Davis and Goadrich, 2006). In binary decisions, confusion matrices measure two independent quantities. The first quantity of interest is the probability of a false alarm. In our example, false alarm probability characterizes how likely it is that a doctor will regard a benign growth as a cancerous mole. The second quantity of interest, known as the true positive rate, will specify the fraction of all cancerous moles that our doctor was able to correctly detect. True positives and false alarms are often examined from the point of view of individual decision-makers. Knowledge of these quantities allows agents to trade off different kinds of errors (Green and Swets, 1988). The notions of false alarms and true positives also turn out to facilitate the development of methods for group decision-making as we will show below in our discussion of collective threat detection (Wolf et al., 2013).

The methodology is applicable to continuous variables as well. As an illustration, we might during some point in the day ask random people on the street to estimate the time of day without looking at the watch and then graph the distribution of opinions to characterize the reliability of their time estimates under our experimental conditions. It is typically useful to have some idea of the reliability of our information sources because the knowledge enables us to estimate the average quality of our final decision, calculate the probability of a serious error or potentially rank different sources in terms of reliability so as to prioritize more reliable sources over less reliable ones (Green and Swets, 1988; Tawn, 1988; Silver, 2012; Marshall et al., 2017). Even more interestingly, it allows us to correct for systematic statistical biases (Geman et al., 2008; Trimmer et al., 2011; Whalen and Yeung, 2015) and, thus, improve overall performance. Systematic biases, if they are measurable, are often easily eliminated by a small change in the decision rule, perhaps similar to how a man who is consistently wrong is easily transformed to a useful assistant if one always acts opposite to his advice. We invite the reader to look at **Figures 1A–C** for a graphical illustration of these issues.

Just like opinions carry information regarding the truth, they may also carry information about each other. As an everyday example, let us look at a group of school children who have been taught to eat or avoid certain types of mushrooms from a common textbook. Our scenario creates an interesting situation, where one need not poll the entire class to know what all kids think. Asking only a few students for their opinion on any particular mushroom will tell us what the others likely think. Their opinions are now generated through a shared underlying mechanism (Barkow et al., 1995) and may be said to have a mutual dependence.

Mutual dependencies between variables influence the optimal decision rule in many ways. Pairs of variables that show a mutual relation to each other are frequently studied using their correlations (though there are other forms of dependencies not captured by correlations). Correlations can impede the emergence of collective intelligence (Bang and Frith, 2017). Thus, in the social sciences, much effort has been devoted to methodologies aimed at eliminating correlations and encouraging independence (Janis, 1972; Myers and Lamm, 1976; Kahneman, 2011), but we will review situations where correlations boost group performance as well. Interestingly, while it is true that if we are using an optimal decision rule, then on average, more information can only improve our performance or leave it at the same level, this conclusion does not hold for suboptimal decision rules. In such cases, extra information can actually decrease the performance (see section 4). Therefore, correlations and dependencies within opinion pools are well worth studying. **Figures 1D–F** illustrates the diverse forms which inter-individual opinion dependencies may take.
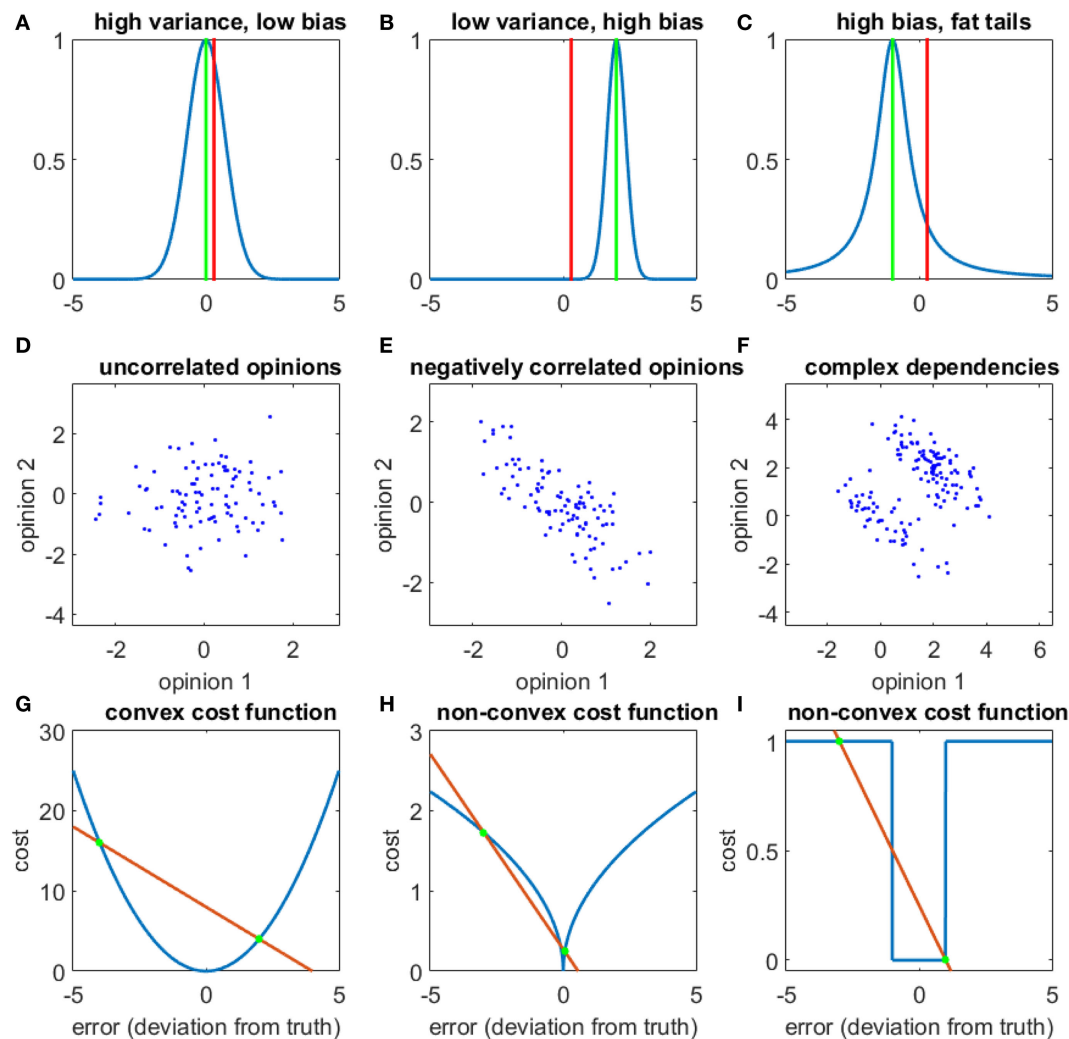
FIGURE 1 | The three factors that influence aggregation rules. **(A–C)**: the relationship between individual opinions and the truth. Blue curves show group opinion distributions, red lines mark the location of the truth, and green lines mark the average group opinion. **(A)** A low bias but high variance distribution. **(B)** A high bias, low variance distribution. **(C)** A biased distribution with fat tails marked by the slower decay of the probability distribution away from the mean. **(D–F)** The relationship between the opinions of two individuals. **(D)** Uncorrelated opinions. **(E)** Negatively correlated opinions. **(F)** A complex dependence between two individual's opinions. **(G–I)** Various cost functions. **(G)** A convex cost function (see Section 5 and Appendix A1.1–1.2 for more extended discussions). **(H,I)** Two non-convex cost functions. In order to illustrate the property of convexity, we have also intersected each cost function with a red line. See section 5 for a further explanation.

After we have determined the relationship between the truth and our information sources as well as the information that the opinions provide about each other, we have all the necessary knowledge to calculate the probability distribution of the truth. Yet knowing the likely values of the truth alone will not be sufficient. Before we are able to produce a final estimate, we need to consider the cost of errors (Green and Swets, 1988). We need a mathematical rule specifying how much cost is incurred by all the various different deviations from the truth which may occur when we make an error. A more extended definition and discussion of cost functions will follow in section 5. At this point, the reader might gain a quick intuition into the topic by examining graphical illustrations of various cost functions in **Figures 1G,H**.

Cost functions are typically application dependent, but in academic papers, the most commonly used cost functions seem to

be the mean squared error and the mean absolute deviation. The cost function has an important influence on the final decision rule. For example, if errors are penalized according to their absolute value, then an optimal expected outcome is achieved if we give as our answer the median of our probability distribution, whereas in the case of the squared error cost function, we should produce the mean of our probability distribution as final answer (Bishop, 2006). As the cost function changes, so changes our decision rule as well. It will turn out that certain cost functions will lead us away from averaging methodologies toward very different decision rules.

Throughout the review, we will make references to the aforementioned three concepts of statistical decision theory and how they have informed the design of new methods for knowledge aggregation. To help structure our review, we have grouped

together methods into subsections according to which factor is most relevant for understanding the aggregation methods, but since ideally an aggregation procedure will make use of all three concepts, a strict separation has not been maintained and all concepts will be relevant to some degree in all subsequent chapters.

3. THE RELATIONSHIP BETWEEN INDIVIDUAL OPINIONS AND THE TRUTH

In decentralized systems, individual agents may possess valuable information about many different aspects of the environment. Ants or bees know the locations of most promising food sources, humans know facts of history, and robots know how to solve certain tasks. But the knowledge of individual agents is usually imperfect to some degree. For the purposes of decision-making and data aggregation, it is useful to have some kind of quantitative characterization of the knowledge of individual agents. Probability distributions and empirical histograms (Rudemo, 1982) are a convenient means to characterize the expected knowledge possessed by a randomly selected individual.

If the truth is known and we have a way to systematically elicit the opinions of random members in a population, then constructing opinion histograms is technically straightforward. Three key characteristics of the empirical histogram are known to be very important for data aggregation: the bias, the variance, and the shape of the distribution (Geman et al., 2008; Hong and Page, 2008). The bias measures the difference between the average group opinion and the truth. The smaller the bias, the more accurate is the group. The variance characterizes the spread of values within the group. If group member opinions have large variance, then we need to poll many people before we gain a good measure of the average group opinion (see Appendix A1.3 for formal mathematical definitions of above terms and **Figures 1A–C** for a pictorial explanation).

The shape of the distribution is a more complicated concept. Many empirical distributions do not have a shape that is easily characterized in words or compact algebraic expressions. If one is lucky enough to find a compact characterization of the distribution it can greatly improve the practical performance of wisdom of the crowd methods (Lorenz et al., 2011; Madirolas and de Polavieja, 2015). In the absence of an explicit description of the distribution, it is helpful to look at qualitative features such as the presence or absence of fat tails. Distributions with fat tails show strong deviations from the Gaussian distribution and are distinguished by unusually frequent observation of very large outliers (Taleb, 2013).

3.1. Leveraging Information about Biases and Shapes

Each of the abovementioned features of the empirical distribution can be leveraged to improve group intelligence. We begin with biases. Biases on individual questions are not very helpful *per se*. When those same biases reliably recur across questions, they become useful. The minds of humans and animals make systematic errors of estimation and decision-making which ultimately stem from our sensory and cognitive architecture (Tversky and Kahneman, 1974; Barkow et al., 1995). These biases can also affect

crowd estimates (Simmons et al., 2011). Whalen made use of the concept of biases for improving crowd estimates of expected movie gross revenues (Whalen and Yeung, 2015). Whalen began by asking people to forecast the gross revenues of various movies. When he graphed crowd averages against the truth an orderly pattern became apparent. Crowds systematically underestimated the revenues of all movies. The bias even appeared greater for higher grossing movies. The remedy to the problem was straightforward—crowd estimates needed to be adjusted to higher values. The up-weighting procedure considerably increased crowd accuracy on a set of hold-out questions which were not used to estimate the bias.

Another important practical use case of biases concerns crowd forecasting of probability distributions. Humans systemically underestimate the probability of high probability events as well as overestimating the probability of low probability events (Kahneman and Tversky, 1979). Human crowd predictions show similar biases and a debiasing transformation can then be used to improve the accuracy of crowd probability predictions (Ungar et al., 2012). A related method uses opinion trimming to improve the calibration of probability forecasts (Jose et al., 2014).

Similar to the way knowledge about biases helps design better aggregation methods, knowledge about the shape of the distribution is critical for designing the best knowledge integration techniques. In many real datasets, varying expertise levels deform the distribution of opinions from a normal distribution to a fat-tailed distribution (Galton, 1907; Yaniv and Milyavsky, 2007; Lorenz et al., 2011). Fat-tailed distributions generate more frequent outliers that have large effects on estimating the mean when using classical statistical procedures. When data are generated from a fat-tailed process, it is better to use robust statistical estimation methods. A useful technique for estimating the mean involves leaving out a certain percentage of the most extreme observations (Rothenberg et al., 1964). Pruning away the outliers may improve wisdom of the crowd estimates (Yaniv and Milyavsky, 2007; Jose and Winkler, 2008). One particular type of distribution called the log-normal distribution even has a convenient estimator known as the geometric mean which can be very effective as an estimation procedure for datasets conforming to the distribution.

3.2. Individuality and Expertise

Previously, we treated all members of the crowd as identical information carriers. This is generally not the case. Sources of information may be distinguished from one another by their type, historical accuracy, or some other characteristic. When information is available regarding the reliability of sources, a weighted arithmetic mean typically works better than simple averaging (Silver, 2012; Budescu and Chen, 2015; Marshall et al., 2017). For example, sites aggregating independent polls produce their final predictions by weighting the independent polls proportionally to the number of participants in each poll, because, all other things being equal, larger polls are more reliable (Silver, 2012).

In the field of multi-agent intelligence, individuals are typically broadly similar, but may nevertheless have some individual characteristics. One particularly frequently explored topic concerns analysis of historical accuracy in order to improve future predictive power. Historical track records are, for example, used to form

smaller but better informed subgroups. Having a subgroup rather than a single expert allows the averaging property to stabilize group estimates whilst avoiding the systematic biases which often plague amateur opinions. Mannes et al. (2014) have studied the performance of select crowds of experts on an extensive collection of 50 datasets. Experts were first ranked relative to past performance and subsequently, the future predictions of either the whole crowd, the best member of the crowd or a collection of the best 5 members of the crowd (the select crowd) were compared with each other. The select crowd method systematically outperformed other methods of knowledge aggregation.

In another study (Goldstein et al., 2014), nearly 100,000 thousand online fantasy football players were ranked in order of past performance. The investigators then formed virtual random subgroups which varied in size and the amount of experts they contained. The behavior of the subgroups was used to predict which players will perform best in English Premier League games. Analysis indicated that small groups of 10–100 top performers clearly out-competed larger crowds where expert influence was diluted, thus showing the benefits of taking expertise into account. In general, following the experts is expected to be beneficial if we have both good track records and there is a wide dispersion in individual competence levels, while for relatively uniform crowds averaging methods perform as well or better (Katsikopoulos and King, 2010).

When extensive historical records are missing, experimental manipulations have been invented to tease out the presence of expertise. One such strategy is known as the wisdom of the resistant (Madirolas and de Polavieja, 2015). Wisdom of the resistant exploits humans' tendency to shift their opinion in response to social information if there is private uncertainty. The natural expectation is for people with more accurate information to have less private uncertainty and to be more resistant to social influence. Wisdom of the resistant methodology consists of a two-part procedure which takes advantage of this hypothesis about human microbehavior. In the protocol, people's private opinions are elicited first and they are subsequently provided with social information in the form of a list of guesses or their mean from other participants to observe how subjects shift their opinion in response to new information. Subjects are ranked in order of increasing social responsiveness and a subgroup with the least flexible opinions is used to calculate a new estimate for the quantity of interest (the exact size of the subgroup is calculated using a p-value based statistical technique so as to still make as much use of the power of averaging as possible). In line with theoretical expectations, the new estimate often improves relative to the wisdom of the crowd (Madirolas and de Polavieja, 2015).

Interestingly, several popular models of decentralized collective movement and decision-making use rules which spontaneously allow the more socially intransigent individuals to have a disproportionately large effect on aggregate group decisions (Couzin et al., 2005; Becker et al., 2017). Natural collectives might, thus, implicitly make use of similar methodologies, although the computation implemented by local rules oriented algorithms is more context dependent (Couzin et al., 2011).

A methodology similar to wisdom of the resistant was recently proposed (Prelec et al., 2017), which asked subjects to predict both

the correct answer and the answer given by the majority. The final group decision was produced by selecting an answer which proved surprisingly popular (more people chose this answer than was predicted by the crowd). Both methodologies leverage the presence of an informed subgroup in the collective and they provide means of identifying informed subgroups without historical track records.

Many of the problems where crowd wisdom is most needed concern areas where there are no known benchmarks or measures of ground truth against which expertise could be evaluated. Under such conditions, we can still determine individual expertise levels by as light reformulation of the problem. Instead of finding the answer to a single question, we again seek to answers to an ensemble of questions. For question ensembles, recent advances in machine learning can be brought to bear on the problem of jointly estimating which answers are correct and who among the crowd are likely to be the experts (Raykar et al., 2010).

As an example, consider the case of a crowd IQ test (Bachrach et al., 2012), where many people fill out the same IQ test in parallel. Here, a machine learning method known as a graphical model is applied to the problem of collective decision-making. The IQ test was an ensemble of 50 questions and IQ was linearly related to the number of correct answers given by the decision-maker (the IQ ranges measured on the test were from 60 to 140). Since individual IQ varies, we can characterize each person with the probability of correctly answering a randomly chosen question on the test, p. We cannot measure p directly, but since the average probability of a correct decision is 75% and the crowd majority will answer most questions correctly most of the time, then we can get an estimate of p by looking at how well each persons answers correlate with the majority vote. These estimated p values can subsequently be used to refine our estimates of which answers are correct, which in turn can be used to refine our estimated p values further. Stepping through this iteration multiple times allows the algorithm to improve on the results of the majority vote.

In the case of crowd IQ, a majority vote among 15 participants produces an average crowd IQ of approximately 115 points, while the machine learning algorithm can be used to boost this performance by a further 2–3 points. It is also interesting to see that unlike what would be expected from Condorcet, crowd IQ effectively plateaus after a group size of 30 is reached. A crowd of 100 individuals has a joint IQ score of merely 120. Given that a group of 100 individuals is very likely to contain a few people with near-genius level (>135) IQ, the study also illustrates why it could sometimes be well worth the effort to find an actual expert rather than relying on the crowd.

Is it possible to utilize expertise if we poll the crowd on a single question rather than on an ensemble? Empirical studies thus far seem to be lacking. We have built a scenario that shows the possibility of improving on the majority vote under some special conditions.

Consider again a crowd of people choosing among some options, where a fraction $1 - k$ will choose their answer randomly, while a fraction of k experts know the correct answer. During actual voting, we sample randomly N individuals from our very large crowd and let them vote. If our crowd members face a choice between two alternatives, then a random member of the crowd

will be correct with probability $p = k + \frac{1}{2}(1 - k)$, and Condorcet theorem will exactly describe how the crowd performance varies as a function of N and p. Suppose we now expand the two-way choice between the correct and incorrect alternative into an K-way choice between the two original choices and $K - 2$ irrelevant distractions. The final opinion is now chosen via a majority vote between the two relevant alternatives while ignoring all opinions landing on the distractions. In the appendix, we prove why the performance of our method for $K > 2$ is always strictly better than the performance of traditional majority voting where the crowd chooses only between 2 alternatives. As is apparent in **Figure 2**, the improvements in performance are quite dramatic, particularly for larger values of K and N. For very large values of K, the performance of the method tends to the same formula as the many-eyes model discussed in the next section (see Appendix A3 for proof).

The efficacy of this hypothetical procedure depends on how closely our assumptions of human micro behavior match with our model. This example merely illustrates that scenarios might be constructed and empirically tested for specific problems which allow investigators to significantly improve performance relative to the Condorcet procedure. Perhaps the closest practical analog to this idea is the use of trap questions in crowd sourcing to filter out people who are insufficiently attentive to their task (Eickhoff and De Vries, 2011).

4. THE ROLE OF DEPENDENCIES

Before we dive into the most catastrophic failures of averaging, it is instructive to once more consider why averaging sometimes works very well. As described above, the majority vote was first analyzed

FIGURE 2 | Using irrelevant alternatives to improve group performance- a simulation study. Group performance curves (calculated from a computer simulation) as a function of the number of total alternatives K using our new voting procedure. The percentage of experts in the crowd is fixed at 10% for this plot. The different colors of curves illustrate how varying the group size N influences group performance for a fixed K. The green curve gives a comparison with Condorcet theorem (which is technically equivalent to the case $N = 2$). See Appendix A3 for proof of why performance always exceeds the Condorcet scenario.

in 18th century France, where Marquis de Condorcet proved his famous theorem demonstrating the efficiency of majority voting for groups composed of independent members (see Appendix A1.4 for a mathematical description of Condorcet voting). A crucial tenet underlying his theorem concerns the assumption of independence (Condorcet, 1785; Boland, 1989; Sumpter, 2010). Condorcet theorem requires more than just a group of individuals who do not interact or influence each other in a social way. It requires the jury members to be statistically independent. In a group with statistically independent members, the vote of any member on a particular issue does not carry any information about how other members of the group voted. For the particular case of Condorcet, if an individual has an expected probability p of producing the correct answer, then we do not need to modify our estimate of the value of p after we learn whether his partner voted correctly or incorrectly.

In all the examples covered in the current section, the aforementioned statistical independence property no longer holds and learning any individual's opinion now also requires us to modify our estimate of his partners' opinions. The lack of statistical independence is not just a feature of our examples. Statistical independence is difficult to guarantee in a species where most individuals have a partially shared cultural background and all members have a shared evolutionary background which constrains how our senses and minds function (Barkow et al., 1995). Because of that shared background, the opinions of non-interacting people are also likely to be correlated in complex ways.

It is easy to notice some ways in which correlations retard collective intelligence. Using the abovementioned example of school children who all learned about mushrooms from a common textbook, we can conclude that in such a scenario, the group essentially behaves as a single person and no independent cancelation of errors takes place (Bang and Frith, 2017). But the influence of opinion dependencies is sometimes even more destructive. We can imagine a group composed of a very large number of members who need to answer a series of questions. On any random question, the probability of receiving a correct answer from a randomly chosen group member is p. Similar to Kuncheva et al. (2003), we can ask what is the worst possible performance of a group with such properties. In the worst-case scenario, questions come in two varieties: easy questions, where all group members know the correct answer, and hard questions, where infinitesimally less than 50% of the people know the correct answer. On the easy questions, the majority vote will lead to a correct answer, while on the hard questions, the majority vote will lead to incorrect decisions. Intuitively, the 50–50 split on the hard questions will ensure that the greatest possible number of correct votes will go to waste since for those questions the correct individual votes do not actually help the group's performance. With such a split of votes, the group will perform as poorly as possible for a given individual level performance (see Kuncheva et al. (2003) for more details). In order for the average person to have an accuracy of p, the proportion of hard questions (t) must satisfy $p = \frac{1}{2}(1 - t) + t$ which means that the group as a whole will be correct in only $2p - 1$ fraction of cases. The result is quite surprising—a group where the average individual is correct 75% of the times may as a whole be correct in only 50% of the questions.

Dependencies, however, are not necessarily detrimental to performance. As we explain in the following two subsections, whether or not correlations and dependencies help or hurt performance depends on the problem at hand (Averbeck et al., 2006; Davis-Stober et al., 2014) and, crucially, on the decision rule used to process the available data. These general conclusions extend to the domain of collective decision-making as well.

4.1. Correlations Can Improve Performance in Voting Models

We begin our discussion of alternative voting procedures with the important example of collective threat detection. Here, the majority vote is eschewed in favor of a different decision rule. A single escape response in a school of fish (Rosenthal et al., 2015) or some-one yelling fire in a crowded room can transition the whole collective into an escape response. A collective escape response begins even though the senses of a vast majority detect nothing wrong with their surroundings. The ability of an individual to trigger a panic is treated very seriously. In the US legal system, one of the few instructions which restricts freedom of speech concerns the prohibition against falsely yelling fire in a crowded room.

Despite the slightly negative connotation of the word, panics are a useful and adaptive phenomenon. For example, panics help herding animals avoid predators after collective detection of a predator (Boland, 2003). Improved collective predator detection and evasion is known as the many-eyes hypothesis and it is thought to be one of the main drivers behind the evolution of cooperative group behavior (Roberts, 1996).

Why is it rational to ignore the many in favor of the few? Consider a very simple probabilistic model to explain this behavior. Let us think of a single agent as a probabilistic detector. Let us also assume that the probability of the agent detecting a predator where none is present is zero, in other words, there are no false alarms. The probability of an animal detecting a predator when one is in fact present is t. The value of t might be much less than one, because detecting an approaching predator is hard unless you happen to catch it in motion or look directly at it. Under the conditions of our scenario, it is clear that other animals will begin an escape only if a predator is in fact present. It follows that if others are escaping, you should begin an escape as well.

For the sake of giving a concrete example (a formal mathematical treatment and derivation of all the formulas related to the panic models which follow are found in the Appendix), let us analyze the case where the probability of a predator attacking is 50 h case makes up 50% of the total incidents. For $t = 0.4$, it gives a value of $p = 0.7$ (**Figure 3**, value at group size = 1). For group sizes larger than 1, the majority vote performs worse than this value because if a predator is indeed attacking, only a minority of the animals will detect the predator and the majority votes that there is no predator present (**Figure 3**, black line). The majority of a large group is then only correct in the 50% of the cases in which there is no attack (**Figure 3**, black curve for large groups).

In real collectives, the majority vote is rejected as a decision rule, and even a single detection by a single member is enough to alert the whole group to the danger. Under such a strategy, the probability of a group correctly detecting a predator increases very rapidly as the group size N increases as $1 - 0.5(1 - t)^N$ (**Figure 3**,
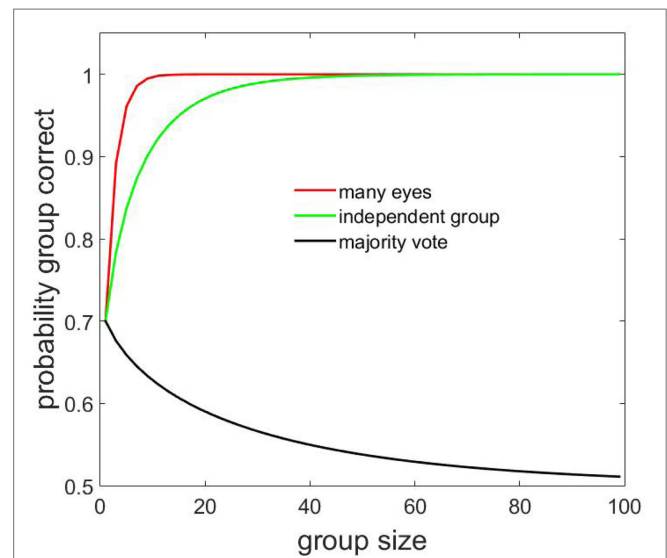
FIGURE 3 | Following the informed minority vs. majority voting. The red curve plots the percent of correct threat assessment as a function of group size for the optimal detection (many eyes) model, where even detection by a single individual can trigger a collective response. The black curve illustrates how performance would change with group size if animals used the majority vote. The green curve plots the performance of a crowd of independent individuals with same individual competence as for the many-eyes model. See main text for details and Appendix A2 for the mathematical derivation of the three lines.

red curve; see Ward et al. (2011) for use of the same expression, known as many-eyes model). The group then detects a predator much more efficiently than if it was relying on the majority.

What would the performance of the group be like if the animals were all statistically independent from each other while retaining the same average individual performance as in the many-eyes model. Now, instead of analyzing the majority decision for the no attack and attack cases separately, we simply plug the probability $p = 0.7$ into the Condorcet majority formula (**Figure 3**, green curve). The plot clearly shows the superiority of the many-eyes model over the independent group. Inter-animal dependencies have increased group intelligence.

The idea of harnessing correlations to increase group performance appears rarely discussed in the voting literature. For example, a recent comprehensive review of group decision-making and cognitive biases in humans had an extensive discussion of how inter-individual correlations can hurt group performance and the ways in which encouraging diversity helps overcome some of the problems (Bang and Frith, 2017). Yet the positive side of correlations and how they may help performance was not covered. Likewise, in another paper on fish decision-making, quorum decision rules were compared against Condorcet's rule as if it was the optimal possible decision rule (Sumpter et al., 2008), even though other rules which account for potential correlations are capable of producing better group performance.

It is also important to note that in the case of applying the majority vote to estimate the presence of a threat, none of the reasons usually provided to explain away the failures of the majority vote apply (Surowiecki, 2004; Kahneman, 2011; Bang and Frith, 2017). The initial votes could be cast completely independently

(without social interaction) and each new vote could add diverse and valuable new information to the pool of knowledge and yet the majority vote would still fail. The insight here is that the majority vote is inappropriate because it does not match with the distribution of knowledge across the collective: a minority has the relevant information about the presence of a predator.

If we examine **Figure 3** more carefully, we see a region corresponding to large group sizes ($N > 45$), where the majority vote for the independent group and the panic model both give near-perfect performance (though the panic model always strictly outperforms the independence model for all $N > 1$, see Appendix A3 for proof). It is, therefore, natural to wonder whether encouraging independence might be a useful practical rule of thumb if one is sure to be dealing with very large groups.

The independence-focused line of reasoning runs into difficulty when one considers the costs necessary to make animals in the group perfectly independent. To guarantee statistical independence, it is not sufficient to merely make the animals in our group weakly interacting. The correlations originate because all animals experience threat or safety simultaneously. Correlations only disappear when the probability of any individual making a mistake is equal in both the threat and the no threat scenario. Any time, the above condition fails to hold, correlations appear, which makes it clear why establishing perfect independence is a precarious task likely to fail in the complexity of the real world. By contrast, the panic model is a robust decision rule, stable against variations in probabilities and guaranteed to give a better than independent performance for small group sizes. It, thus, becomes more apparent why natural systems have preferred to adapt to and even encourage correlations rather than fight to establish independence.

We note that even for the many-eyes model, in practice there is usually a small probability of a false alarm, and field evidence from ornithology demonstrates how animals can compensate against rising false alarm rates by raising the threshold for the minimum number of responding individuals necessary to trigger a panic (Lima, 1995). We point the interested reader to Appendix A2 for a mathematical treatment of the false alarm scenario.

The decision rule adopted by vigilant prey could be called a "full vote." In order to declare a situation safe, all individuals must agree with the proposition. A similar rule has been rediscovered in medical diagnostics. In medical diagnostics, some symptoms such as chest pain are inherently ambiguous. Sudden chest pain could signal quite a few possible conditions such as a heart attack, acid reflux, a panic attack, or indigestion. In order to declare a patient healthy, she must pass under the care of a cardiologist, a gastro-enterologist, and a mental health professional. All experts must declare a patient healthy before he can be released from an examination. In the case of a panel of experts, their non-overlapping domains of expertise help insure the effectiveness of the full vote.

A similar idea has been implemented in the context of using artificial neural networks (a machine learning method, see Section 6 for more details) to detect lung cancer in images of histological sections (Zhou et al., 2002). An ensemble of detectors is trained using a modified cost function which heavily penalizes individual neural networks when they declare a section falsely malignant.

The training procedure makes false alarms rare, so the full vote procedure can be used to detect cancer more efficiently than if the networks had been stimulated to be maximally independent.

4.2. Correlations and Continuous Variables

In the case of averaging opinions about a continuous quantity, correlations also have a profound effect on group performance. The average error on a continuous averaging task is given by the sum of the bias and the variance (Hong and Page, 2008). Variance declines as we average the opinions of progressively larger pools of opinions (Mannes, 2009). Correlations control how rapidly the variance diminishes with group size. The speed of decrease is slowest when correlations are positive. Finding conditions where errors are independent helps speed up the decrease of variance. The most rapid decrease occurs when correlations are negative (Davis-Stober et al., 2014). For large negative correlations, the errors in pairs of individuals almost exactly cancel and even a very small group can function as well as a large crowd of independent individuals. The benefits of negative correlations are exploited in a machine learning technique termed negative correlation learning (Liu and Yao, 1999).

Correlations can be leveraged most efficiently when we have individual historical data. Personalized historical records enable the researcher to estimate separate correlation coefficients for every pair and compute the optimal weighting for every individual opinion. The benefits of correlation-based weighting are routinely applied in neural decoding procedures, where the crowd is composed of groups neurons and opinions are replaced by measurements of neural activity. Averbeck et al. (2006), for example, study the errors induced in decoding if neural activity correlations are ignored, and find that ignoring correlations generally decreases the performance of decoders when compared to the optimal decoder which takes the information present in correlations into account. Similar to Davis-Stober et al. (2014) who study correlated opinions, they find a range of situations where correlations improve decoding accuracy as compared to independently activating neurons.

5. THE ROLE OF COST FUNCTIONS

5.1. Measures of Intelligence

Collective intelligence is of course a partly empirical subject. After the theoretical work of Condorcet, the next seminal work in the academic history of wisdom of crowds comes from Galton, whose work we briefly described in the introduction. The conclusion of his study was that simple averaging of individual estimates is, as an empirical matter, a more useful way to estimate quantities than relying on faulty individual opinions. In addition to Galton's work, another classic study of crowd intelligence involved subjects estimating the number of jelly beans or marbles contained in a jar (Treynor, 1987; Krause et al., 2011; King et al., 2012). The true number of beans is typically between 500 and 1,000, so exact counting is not feasible for the subjects. If the crowd is larger than 50 individuals, the crowd median and/or mean opinion typically comes within a few percent of the true value. The effect is even somewhat independent of the sensory modality involved. In a study of somatosensory perception, 56 children estimated the

temperature of their class room. The average of their 56 guesses deviated from the true value by just 0.4°(Lorge et al., 1958).

Galton and many others who followed gave empirical demonstrations regarding the remarkable effectiveness of simple averaging without any mathematical arguments as to why the phenomenon occurs. Perhaps because the performance of the crowd in these early studies was spectacularly good, there was also a lack of explicit comparison to other ways of making decisions. In more recent years, there has been more focus on the failure of crowds. Many examples are known where crowds fail to come close to the truth (Lorenz et al., 2011; Simmons et al., 2011; Whalen and Yeung, 2015). Lorenz et al. (2011) report an average crowd error of nearly 60% (relative to the truth) in a set of tasks consisting of estimating various geographical and demographic facts. In psychology, there is a rich literature on the heuristics and biases utilized in human decision-making (Tversky and Kahneman, 1974), which can also bias crowd estimates (Simmons et al., 2011).

Examining collective performance in cases where the crowd makes practically significant mistakes led to a need to perform more explicit comparisons between different methodologies. It is common to compare wisdom of crowd estimates with the choosing strategy.

In the choosing strategy, we pick one opinion from the crowd at random and use that opinion as our final estimate. To quantitatively compare averaging and choosing, we first measure the error of a guess as error = |our guess − true value|, where |x| stands for absolute value of any number x. To assess the impact of an error, we also have to specify a cost function. A cost function is a mathematical measure which specifies how damaging an error is to overall performance. The smaller the overall cost, the better the performance. Common cost functions found in the literature are the absolute error (also called the mean absolute deviation) and the squared error cost functions. If we are using the choosing strategy, then the error will typically be highly variable from person to person, because individual guesses are variable. In order to compare the performance of the choosing strategy with the performance of the crowd average opinion, we average the costs of individual guesses and then compare the average cost with the cost of the mean crowd opinion.

We illustrate the role of a cost function with a numerical example. In an imaginary poll, we query four people about the height of a person whose true height is 180 cm. The group provides four estimates: 178, 180, 182, and 192 cm. The corresponding error values are |178 − 180| = 2, |180 − 180| = 0, |182 − 180| = 2, and |192 − 180| = 12. The mean absolute deviation cost is (2 + 0 + 2 + 12)/4 = 4. Since the crowd mean is 183, the crowd opinion induces a cost of 3 only. In this example, averaging outperformed choosing. Similarly, for the squared error cost function, the choosing strategy has an expected error of $(2^2 + 0^2 + 2^2 + 12^2)/4 = 37$, while the crowd mean causes an error of $(183 − 180)^2 = 9$. The crowd mean again outperforms random choice.

It has become common practice to emphasize the superiority of wisdom of crowd estimates over the choosing strategy with performance measured through use of the mean absolute deviation or the mean squared error cost function (Hong and Page, 2008; Soll

and Larrick, 2009; Manski, 2016). An unconscious reason behind the popularity of the comparison might be that it will always yield a result that casts collective wisdom in a favorable light. A mathematical theorem known as Jensen's inequality guarantees the superiority of the average over the choosing strategy for all convex cost functions. The mean squared error and the mean absolute deviation are both examples of convex cost functions.

The exact definition of a convex function is rather technical (see Appendix A1.1–1.2 for a formal definition of both convexity and Jensen's inequality), but we may gain some intuition into the concept if we examine what happens if we intersect various cost functions in **Figures 1G–I** with randomly drawn lines. For each panel, if we focus on the relationship between the red line and the blue curve in between the green dots, we see that for panel G the red line is always above the blue curve, whereas for H and I, the red line may be either above or below the blue line depending on which region between the green dots we focus on. In fact, for function G, the blue curve is always below the red line for any possible red line we may think of as long as we focus on the region that is between the two points where the particular line and the curve intersect. It is this property that makes G a convex function and allows us to guarantee that the group average error is always smaller than the average individual error.

Some authors have elevated Jensen's inequality and similar mathematical theorems to the status of a principle which justifies the effectiveness of collective intelligence (Surowiecki, 2004; Larrick and Soll, 2006; Hong and Page, 2008). We hold ourselves closer to the position of authors who have questioned these and similar conclusions (Manski, 2016). Fundamentally, Jensen's inequality is merely a property of functions and numbers. We might sample 100 random numbers from a computer and use them to estimate the year Winston Churchill died. If I measure my performance using convex cost functions, then the average of my sample will induce a lower cost than a choosing strategy. Should I say that the collection of random numbers possesses collective intelligence?

Furthermore, reporting collective performance on a single question using a single numerical measure exposes the investigators to an unconscious threat of cherry-picking. Perhaps the good performance of the crowd was simply an accidental coinciding of the crowd opinion with the true value of one of the many possible questions that many investigators have proposed to crowds over the years.

Instead, we advocate the study of correlations on ensembles of questions as was recently also done by Whalen and Yeung (2015). We illustrate the procedure by reanalysis of a dataset from the study by Yaniv and Milyavsky (2007), where students were asked to estimate various historical dates. On **Figure 4**, we have plotted the true values versus the wisdom of crowd estimates for 24 questions. Such an analysis gives a good visual overview of the data. For example, it is immediately clear from the plot that wisdom of crowd estimates are strongly correlated with the truth across the ensemble and there is clearly knowledge present in the collective. We find that on an average question, the crowd wisdom missed the truth by nearly 30 years. On certain questions, the crowd error was undetectable, while on others the crowd was off by nearly 100 years. Overall, the collective performance is of mixed

FIGURE 4 | Wisdom of the crowd for historical dates. Mean of 50 independent opinions versus true values for 24 questions concerning historical dates. The red line marks the ideal performance curve; real performance frequently strongly deviates from that line. Data from Yaniv and Milyavsky (2007).

FIGURE 5 | Comparison of averaging and choosing strategies for different cost functions. Cost incurred by the averaging strategy (red curve) and the choosing strategy (blue curve) as a function of the location of the true value for four different cost functions. The five opinion values on which the performance is calculated: -1, -0.5, 0, 0.5, and 1. Quadratic cost is strictly convex, and mean (red) is then always below choosing (blue). Absolute error cost function is weakly convex, and mean (red) is then always below or equal to choosing (blue). Square the third and fourth curves are neither convex nor concave. The threshold cost function gives a cost of 0 if an opinion is closer than 0.45 to the truth and a cost of 1 for all other values.

quality, with excellent performance on some questions, mediocre performance on others, and no clear systematic biases.

5.2. Beyond Convexity

Aside from the fact that Jensen's inequality may be applied to any collection of numbers, there is another problem with analysis of collective performance as they are currently commonly carried out. There is an exclusive focus on convex cost functions. Yet many real-world cost functions are non-convex. In a history test, problems will typically have only a single acceptable answer. A person who believes the US became independent in 1770 will receive zero points for his reply, just like a person who believes the event took place in 1764, even though the first person is twice as close to the truth. Similarly, an egg which was cooked for 40 min too long is not substantially better than an egg over-cooked for 120 min as both are inedible and should induce similar costs for the cook.

What happens to the performances of the averaging and the choosing strategies when we change our cost from a convex to a non-convex function? We will once again make use of our afore-mentioned example of guessing heights. As our new cost function, we will use a rule which gives a penalty one to all examples that deviate from the truth by more than 1 cm and assigns a cost of 0 to answers which are less than 1 cm away from the truth. Our set of opinions was 178, 180, 182, and 192 cm with the true value lying at 180 cm. In this case, the crowd mean has a penalty of 1, because the crowd mean of 183 misses the true value of 180 by more than 1. Three out of four individual guesses also miss the truth by more than a year, but one guess hits the truth exactly, so the average cost of the choosing strategy is $(1+0+1+1)/4 = 0.75$. In this case, the crowd mean underperforms relative to the choosing strategy. A similar effect results from using a cost function which penalizes guesses according to the square-root of their absolute error. The square-root cost function penalizes larger errors more than smaller errors, but the penalty grows progressively more slowly as

errors increase. The crowd mean has a cost of $\sqrt{183-180} = 1.7$. The choosing strategy has an expected cost of $(\sqrt{2}+\sqrt{0}+\sqrt{2}+\sqrt{12})/4 = 1.6$. The crowd mean incurred a higher expected cost than a randomly chosen opinion. Our examples illustrate that the best strategy for opinion aggregation is highly dependent on the cost function.

A different way to visualize the same result would be to consider the cost incurred by the same pool of opinions as the location of the truth varies. In **Figure 5**, we consider the cost performance of a fixed pool of 5 opinions (with values -1, -0.5, 0, 0.5, and 1) as a function of the location of the true value. As can be seen from the graph, convex cost functions such as the mean square error and the mean absolute deviation produce a lower error when the mean opinion is used independently of where the true value is located. Non-convex functions such as the mean square root of the absolute deviation reveal a more complex picture. Sometimes it is better to choose and sometimes it is better to average. No simple optimal prescription is possible.

If the cost function is not convex, then Jensen's inequality no longer applies and averaging is not guaranteed to outperform choosing. As our last two examples showed, the opposite might be the case. In that light, it is intriguing to note that when humans take advice from other people, they often opt for a choosing strategy rather than an averaging strategy (Soll and Larrick, 2009). This behavior has been seen as suboptimal (Yaniv, 2004; Mannes,

2009; Soll and Larrick, 2009), but it may in fact be a rather rational behavior. The human crowd often contains a substantial fraction of experts who know the answers to certain questions while other members of the crowd have less information about the question at hand. If we assume that advice becomes beneficial only if it reaches relatively close to the truth, then it becomes rational to pick a random opinion in the hopes of hitting expert advice, rather than relying on the crowd mean, which might lie far from the truth because of distortions by non-expert advice.

To analyze the problem more systematically, we have re-examined an experiment from Yaniv and Milyavsky (2007), where 150 students were individually presented with questions about when 24 prominent historical events took place. They were subsequently provided with advice from two, four, or eight other students. The students had the option of combining their initial private opinion with further advice (the advice was anonymous and was not presented in person) from other subjects. They had financial incentives to provide maximally accurate answers in both the individual and the advice-taking part of the experiment. We found that after receiving anonymous advice, in approximately 70% of cases, subjects stayed with their initial private opinion or chose the opinion of one particular adviser as their final answer. In **Figure 6**, we plot the cumulative distribution of errors of the students initial private estimates (red) and their revised opinions after hearing advice (blue) from 2 (left), 4 (middle), or 8 advisers (right). These may be compared against a strategy of averaging the student opinion and all the advisory opinions received (**Figure 6**, black). The distribution of errors indicates that students adopt a strategy that produces a more frequent occurrence of low error answers than the averaging strategy (though, of course, as guaranteed by Jensen's inequality, the mean absolute error of the averaging strategy is lower than the choosing strategy and the strategy adopted by the student population as a whole. The aggregate gains of averaging with respect to squared errors mainly originate from the reduced occurrence of extreme errors in the averaging strategy).

It has been argued that durable real-world systems should evolve to a point where costs must be concave in the region of large errors as a robust design against large outliers (Taleb, 2013). It is interesting to speculate that human advice-taking diverges from the averaging strategy precisely because it takes advantage of non-convexity. So far, advice taking on everyday tasks has been understudied, possibly due to methodological difficulties. In the future, it will be illuminating to compare performance of choosing and averaging strategies on more naturalistic problems.

6. EMBRACING COMPLEXITY: A MACHINE LEARNING APPROACH

Previous research has primarily emphasized how simple rules of opinion aggregation can often produce remarkable gains in accuracy on collective estimation tasks. Yet we have also shown that such simple rules may fail in unexpected ways. We have outlined many possible sources of failure, which tend to occur if any of the following conditions are true:

1. The cost function is not convex.
2. The distribution of knowledge within the collective is inhomogeneous.
3. The pool of crowd opinions is not composed of statistically independent estimates.
4. The distribution of opinions has fat tails.
5. The crowd has significant and systematic biases.

One way to deal with these pitfalls is to use domain knowledge to design new estimation heuristics to compensate for the deficiencies in simpler methods. This approach has been successful and we have given several examples of their utility in practical applications. But these new heuristics often lack the mechanical simplicity of the averaging prescription and risk lacking robustness against unaccounted factors of variation in crowd characteristics.
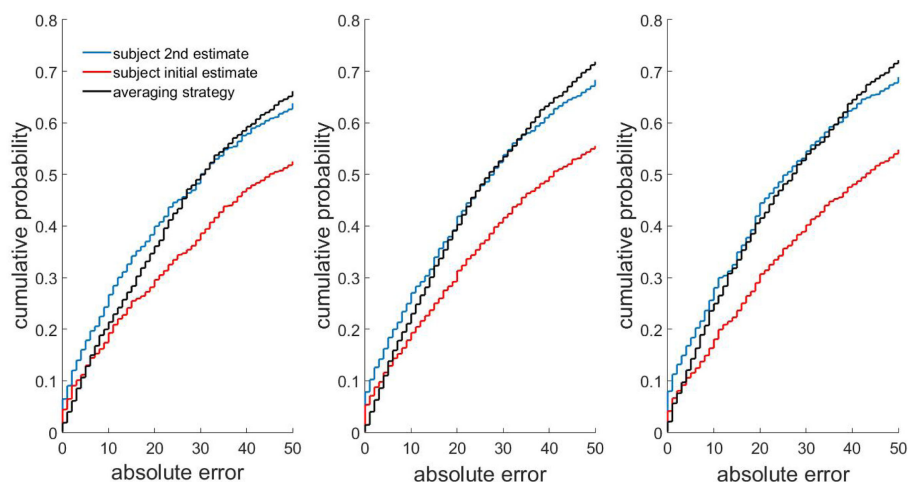
FIGURE 6 | Errors in human advice-taking strategy compared to the averaging strategy. The cumulative distribution of errors of three advice-taking strategies for 2, 4, and 8 advisers. Red curve: initial subject opinion. Blue curve: subject opinions after hearing advice from 2 (left), 4 (middle), or 8 (right) randomly chosen fellows. Black curve: averaging strategy, which calculates a final estimate mechanically by averaging a subjects initial opinion together with all advisers opinions. Data from Yaniv and Milyavsky (2007).

It is the issue of unaccounted characteristics which should be most troubling to the theoreticians. It is easy to perform mathematical analysis of simple models such as Condorcet voting, but as we have previously shown, the confidence derived from such theoretical guarantees has a false allure. Usually, we do not have complete knowledge of all the complex statistical dependencies that occur in the real world and, therefore, the behavior of simple decision rules is liable to unpredictable in practice. The issue of practical unpredictability motivates us to examine ways to create collective intelligence in a way that makes more direct contact with the idea of optimizing real-world performance.

An appealing way to deal with greater complexity is to rely more on methodologies that incorporate complexity into their foundations. Machine learning is capable of learning decision aggregation rules directly from data and can be used to design computational heuristics in a data-driven manner. It can be used to either verify the optimality or near-optimality of known heuristics on a given task or to design new aggregation methods from scratch. While machine learning methods may on occasion be less intuitive for the user, they come with performance guarantees because they are inherently developed by optimizing performance on real-world data. The black box nature is a necessary price which one must pay for the ability to deal with arbitrarily complex dependencies.

Neural networks are one class of machine learning methods that allow the aforementioned procedure to be carried out automatically (LeCun et al., 2015). A neural network is composed of artificial neuron-like elements that transform input opinions into an output estimate. If the researcher has access to a dataset where the true value of the estimated quantity as well as the pool of crowd opinions are known for many groups, then it is possible to find a very close approximation of the optimal decision rule that brings the input opinions into desired outputs.

We will next illustrate the application of neural network-based methods for a simulated dataset, where we find the optimal decision rule in a data-driven manner, and we also apply the method to a cancer dataset where we show that a network has a better performance than the majority vote and previously proposed heuristics.

In our hypothetical example, we consider a group of 30 people that have repeatedly answered questions about historical dates. In our simulated crowd, 50% of individuals will know the answer approximately (their opinion will have a SD of ± 0.1 around the true value). The other 50% are less informed and present a bias to lower values (mean bias -1 ± 0.2). Under this scenario, it is intuitively clear that an optimal decision rule would look for clusters within the pool of opinions and the network must also learn to ignore the opinions coming from the lower cluster. We examined whether a neural network would be able to learn a similar decision rule entirely from data. For our scenario, the crowd mean strategy had an average error of 0.50 whereas a neural network trained on opinion groups was able to reduce the average error to 0.04 (see Appendix A4 for details on training and network architecture), thus demonstrating that neural networks can learn useful approximations to reduce the average error.

We have also examined whether neural networks could improve upon the performance of previously proposed heuristics on a skin

cancer classification dataset (Kurvers et al., 2016). In the dataset, forty doctors had given their estimations and subjective confidence scores (four point scale) on whether particular patients had malignant melanoma by examining images of their skin lesions. As in Kurvers et al. (2016), we used Youden's index as a measure of accuracy, given by $J =$ sensitivity + specificity − 1, with sensitivity defined as the proportion of positive cases correctly evaluated and specificity defined as the proportion of negative cases correctly evaluated. This measure weights equally sensitivity and specificity and it is, thus, insensitive to the unbalances of a dataset (in this case, more cases without cancer than with cancer). We then generated virtual groups of doctors and examined the accuracy of their aggregated judgments. If all doctors in the group agreed on a diagnosis, their joint shared opinion was used as the diagnosis. If there was disagreement, we compared the performance of the following three heuristics for conflict resolution:

1. Use the opinion of the more accurate doctor in the group ("best").
2. Use the opinion of the more confident doctor ("confident").
3. Use the opinion held by the majority ("majority").

In the "best" and "confident" heuristics, if the higher accuracy or confidence was shared by more than one doctor, the majority opinion within that subgroup was selected. If in spite of all the selection rules, there still was a tie, 0.5 was added to the count of correct answers and 0.5 to the mistakes.

We also fitted a neural network that was given as input the historical accuracies of the doctors, their diagnosis on each case, and their declared confidence scores. For any input, the output of the network gave the probability of the given input being consistent with a cancer diagnosis. If the probability exceeded 50%, then the network output was counted as giving a cancer diagnosis. We asked whether a network can find an aggregation decision rule better than the heuristics. The network (a multilayer perceptron) trained with backpropagation on 50% of the data. Another 25% of the data were used as validation dataset. To minimize overfitting, we used the early stopping procedure, where the weights of our network are saved during every epoch of training and in our final testing, we use the version of the weights which gave highest performance on the validation dataset. Testing of the network was done in the remaining 25% of the dataset. **Figure 7A** gives the learning curve of one network on the test dataset depending on the number of training epochs. As an example, for groups of five doctors, we found mean network performance of $J = 0.804$ and SD of 0.060. The different heuristics had the following performance for the same data: 0.757 ± 0.060 ("best"), 0.767 ± 0.067 ("confident"), and 0.801 ± 0.061 ("majority"); see **Figure 7B** for mean improvement of network over heuristics.

For groups of 2, 3, 5, and 7 doctors, we trained 50 networks using different $50 - 25 - 25\%$ partitions of the data into training, validation, and test. We found that both the network and the three heuristics proposed improved their performance over the test cases for increasing group sizes (**Figure 7C**). The networks not only were more accurate than the rest of the heuristics for every group size (except against the majority voting for groups of 3 doctors) but also consistently better in every single partition
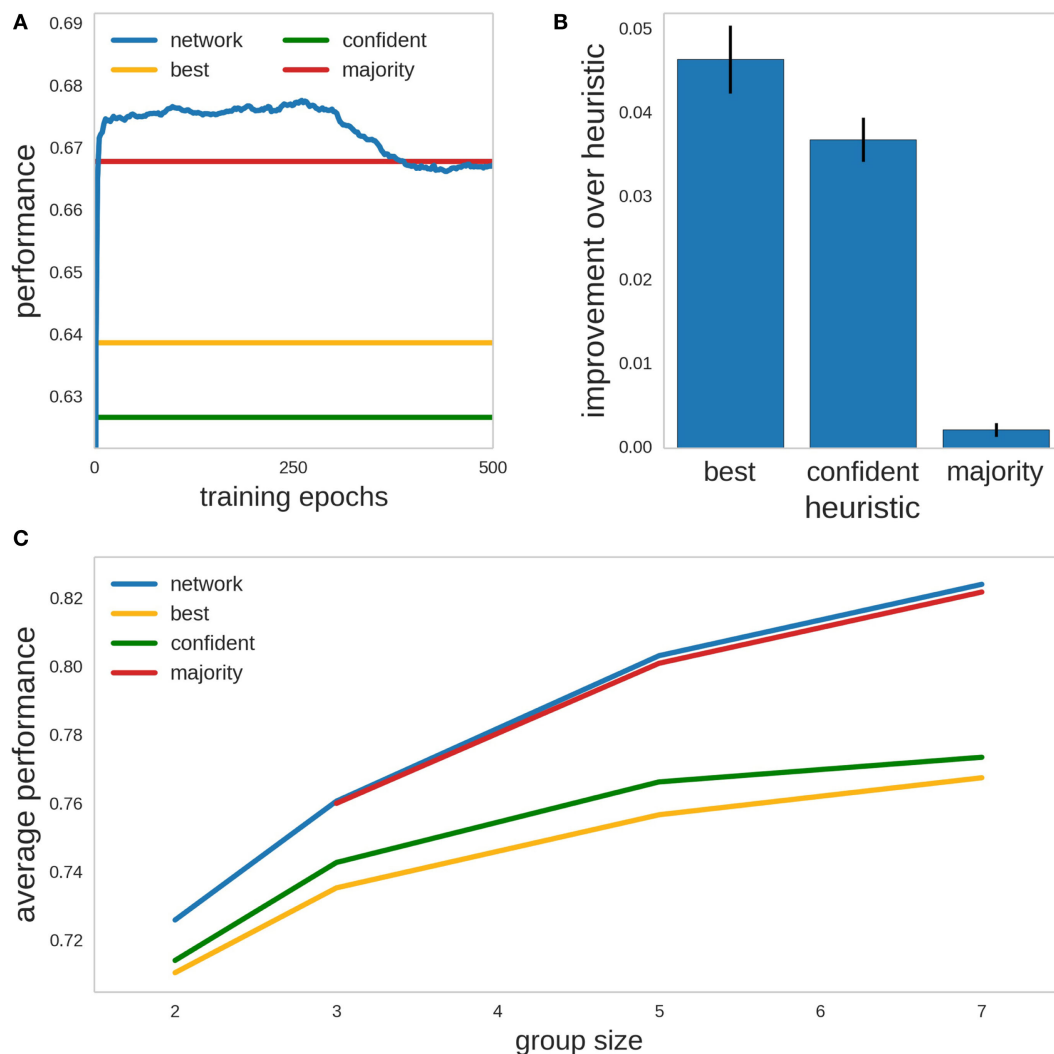
FIGURE 7 | Network learning how to combine the opinions and confidence scores of three doctors into a cancer/no cancer classification rule. **(A)** Example of the training of one network for a particular partition of the data into 50% for training and 25% each for validation and test, for groups of 5 doctors. Shown is the evolution of the performance for the test data (Youden's index, $J =$ sensitivity $+$ specificity $- 1$) of the network (blue line), and the performance of the best doctor heuristic (yellow line), the more confident heuristic (green line), and the majority heuristic (red line). **(B)** Mean improvement in Youden's index of the network over the heuristics for groups of 5 doctors. Error bars are SEM. **(C)** Average performance of network and heuristics over all the validation sets, for groups of 2, 3, 5, and 7 doctors. Colors as in panel **(A)**.

of the cases into training, validation, and test ("best" and "confident": $p < 10^{-5}$ for all group sizes; "majority": $p = 0.0098, 0.027$ for $n = 5, 7$. Wilcoxon signed-rank test). Overall, the difference between the optimal decision rule found by the network and the majority rule is small in this dataset and another way to view the results would be to say that the analysis through use of neural networks gives the user confidence that the majority rule is near-optimal for the present dataset. Note that we were unable to extend our analysis above the case of n = 7, because the permutation procedure we used to create pseudo-groups contains progressively greater overlaps for higher n since we are sampling from a limited pool of 40 doctors, and the statistical independence of our pseudo-groups is no longer guaranteed for larger n, which prevents reliable calculation of p-values.

7. DISCUSSION

The collection of methodologies grouped under the umbrella term wisdom of crowds (WOC) has found widespread application and continues to generate new research at a considerable pace. As the number of real-world domains where WOC methods have been applied increases, researchers are beginning to appreciate that each new domain requires considerable tuning of older methods in order to reach optimal performance. Early focus on universal simple strategies (Condorcet, 1785; Surowiecki, 2004; Hastie and Kameda, 2005) has been replaced with a plethora of methods that have sought to find a better match between the problem and the solution and by doing so have shown increases in performance relative to the averaging baseline (Goldstein et al., 2014; Budescu

and Chen, 2015; Madirolas and de Polavieja, 2015; Whalen and Yeung, 2015).

Many new avenues of research remain to be explored. Machine learning tools and improved ability to gather data provides the opportunity to learn more sophisticated WOC methods in a data-driven fashion (Rokach, 2010; Bachrach et al., 2012; Polikar, 2012; Sun et al., 2017). We are likely to learn much more about effective strategies of opinion aggregation through their widespread adoption. It will also be important to explore whether machine learning rules can be made intelligible to the end user. Techniques such as grammatical evolution (O'Neil and Ryan, 2003), symbolic regression (Schmidt and Lipson, 2009), and the use of neural networks with more constrained architectures may provide a potential approach to the problem.

Hopefully, a synergistic interaction will also continue to take place between the study of collective wisdom and the field of swarm robotics, which seeks to find better ways to coordinate the activities of small independent robots who work together to achieve joint tasks (Bonabeau et al., 1999). One particular area where synergy might be achieved concerns finding a better integration between the methods of task allocation and consensus achievement (Brambilla et al., 2013). It will be interesting to see whether within-swarm task allocation methods could be combined with methods of consensus achievement to simultaneously encourage both diversity of expertise and cooperative action, similar to how crowd intelligence methods benefit from context-dependent reliance on experts (Zhou et al., 2002; Ward et al., 2011; Goldstein et al., 2014). Similar ideas have already borne fruit in the training of expert ensembles of neural networks (Zhou et al., 2002).

Looking toward other unexplored directions, perhaps the least explored avenue in the field of collective wisdom concerns the formulation of the question itself. When human beings describe their choices, they leave a lot of assumptions unstated (Kahneman, 2011). For example, both Alice and Mark may say that they enjoy vacations in France, but once we specify that Mark spent his time in museums while Alice spent her time in the mountains, it becomes obvious that the phrase "liking France" means very different things for the two people. This problem poses a challenge for collective decision-making. Let us suppose that on a scale of 1–10, Mark and Alice give the following rating to his personal experiences: hiking in France is 1(M), 7(A), museums in France 7(M), 1(A) and diving in Egypt: 5(M), 5(A). If Alice and Mark are now planning a joint trip and decide just between going to France or Egypt without being more specific, then they might decide to go to France. After all, if each spent all their time in France engaged in their privately preferred activity, then the expected value of their experience would be 7. But when they actually go to France, they suddenly discover that their preferences conflict and whatever activity they attempt as a group, their average enjoyment of France only has a rating of 4, which would be lower than the joint rating for Egypt.

Note that despite superficial similarity of our example to the problems studied in the literature on social choice (Sumpter, 2010; List, 2011), the dilemma here is in fact caused by an entirely different phenomenon which is more psychological than mathematical in nature. Humans make temporally local decisions based on expected future plans. Sometimes those plans may be implicit rather than explicit, which will lead to hidden conflict even if the two parties appear in agreement at present moment. If the formulation of the alternatives does not take this complexity into account, then we may find ourselves making suboptimal collective decisions. One promising approach, known as the wiki-surveys (Salganik and Levy, 2015), has opened up the formulation of alternatives to the crowd as well. In wiki-surveys, responders are allowed to not just rate alternatives but to provide new alternatives as well, which presumably allows alternatives to be formulated in a more naturalistic manner. On the theoretical side, an integration between the fields of reinforcement learning (which studies temporally extended decision-making) and collective intelligence may provide a fruitful theoretical framework in which to further explore these problems (Biro et al., 2016).

Another potentially exciting and under-explored question concerns research into how and why the distribution of human knowledge comes to have a variety of different classes of distributions in different domains. Related to this, it is crucial to study the shaping of collective knowledge as well. Social and educational policies can presumably direct the distribution and development of human expertise. It could be useful to examine what kind of policies will be most cost-effective in facilitating group intelligence. For some domains, the answer will probably rely on encouraging wide and diverse participation (Page, 2008), whereas for other domains, selective filtering and resource investment into a small group of experts (Goldstein et al., 2014; Budescu and Chen, 2015) might provide a more cost-effective way to increase collective knowledge. As an example, we point the reader to the recent article about reward schemes that encourage holding a correct minority and how such schemes improve collective performance (Mann and Helbing, 2017).

We attempted to demonstrate that far from being a mostly solved problem with well-established standard methodologies, the field of collective intelligence is rather in a state of rapid innovation, with new context-specific heuristics being rapidly developed and many exciting questions remaining under explored. We hope to have shown how to integrate current methodologies into a common framework, which can potentially further stimulate research into the open problems on both the empirical and theoretical sides as well.

AUTHOR CONTRIBUTIONS

AL and GP revised the review; AL, GM, and GP analyzed data; and AL and GP wrote the review.

ACKNOWLEDGMENTS

FUNDING

REFERENCES

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., et al. (2016). Tensorflow: Large-Scale Machine Learning On Heterogeneous Distributed Systems. *arXiv preprint arXiv:1603.04467*.

Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366. doi:10.1038/nrn1888

Bachrach, Y., Graepel, T., Kasneci, G., Kosinski, M., and Gael, J. V. (2012). "Crowd IQ -aggregating opinions to boost performance," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent*, Valencia.

Bang, D., and Frith, C. D. (2017). Making better decisions in groups. *R. Soc. Open Sci.* 4, 170193. doi:10.1098/rsos.170193

Barkow, J. H., Cosmides, L., and Tooby, J. (1995). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. USA: Oxford University Press.

Becker, J., Brackbill, D., and Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proc. Natl. Acad. Sci. U.S.A* 114, E5070–E5076. doi:10.1073/pnas.1615978114

Biro, D., Sasaki, T., and Portugal, S. J. (2016). Bringing a time-depth perspective to collective animal behaviour. *Trends Ecol. Evol.* 31, 550–562. doi:10.1016/j.tree.2016.03.018

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York, NY: Springer.

Boland, C. R. J. (2003). An experimental test of predator detection rates using groups of free-living emus. *Ethology* 109, 209–222. doi:10.1046/j.1439-0310.2003.00860.x

Boland, P. J. (1989). Majority systems and the Condorcet Jury theorem. *Statistician* 38, 181. doi:10.2307/2348873

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. New York: Oxford University Press, 1.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.* 7, 1–41. doi:10.1007/s11721-012-0075-2

Budescu, D. V., and Chen, E. (2015). Identifying expertise to extract the wisdom of crowds. *Manage. Sci.* 61, 267–280. doi:10.1287/mnsc.2014.1909

Clemen, R. T. (1989). Combining forecasts: a review and annotated bibliography. *Int. J. Forecast.* 5, 559–583. doi:10.1016/0169-2070(89)90012-5

Condorcet, M. D. (1785). *Essai sur l'application de l'analyse Ã la probabilitÃ©des decisions rendues Ã la pluralitÃ©des voix*. Paris: Bibloteque Royale.

Conradt, L., and Roper, T. J. (2003). Group decision-making in animals. *Nature* 421, 155–158. doi:10.1038/nature01294

Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., et al. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science* 334, 1578–1580. doi:10.1126/science.1210280

Couzin, I. D., and Krause, J. (2003). Self-organization and collective behavior in vertebrates. *Adv. Study Behav.* 32, 1–75. doi:10.1016/S0065-3454(03)01001-5

Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature* 433, 513. doi:10.1038/nature03236

Dall, S. R., Giraldeau, L.-A., Olsson, O., McNamara, J. M., and Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends Ecol. Evol.* 20, 187–193. doi:10.1016/j.tree.2005.01.010

Davis, J., and Goadrich, M. (2006). "The relationship between precision-recall and roc curves," in *Proceedings of the 23rd International Conference on Machine Learning* (Pittsburgh: ACM), 233–240.

Davis-Stober, C. P., Budescu, D. V., Dana, J., and Broomell, S. B. (2014). When is a crowd wise? *Decision* 1, 79–101. doi:10.1037/dec0000004

Dietterich, T. G. (2000). *Ensemble Methods in Machine Learning*. Berlin, Heidelberg: Springer, 1–15.

Eickhoff, C., and De Vries, A. (2011). How crowdsourcable is your task? in *Proceedings of the workshop on crowdsourcing for search and data mining (CSDM) at the fourth ACM international conference on web search and data mining (WSDM)*, Hong Kong, China.

Galton, F. (1907). Vox Populi (the wisdom of crowds). *Nature* 75, 450–451. doi:10.1038/075450a0

Garnier, S., Gautrais, J., and Theraulaz, G. (2007). The biological principles of swarm intelligence. *Swarm Intell.* 1, 3–31. doi:10.1007/s11721-007-0004-y

Geman, S., Bienenstock, E., and Doursat, R. (2008). Neural networks and the bias/variance dilemma. *Neural Computation* 4, 1–58. doi:10.1162/neco.1992.4.1.1

Goldstein, D. G., McAfee, R. P., and Suri, S. (2014). "The wisdom of smaller, smarter crowds," in *Proceedings of the Fifteenth ACM Conference on Economics and Computation – EC '14* (New York, NY: ACM Press), 471–488.

Green, D. M., and Swets, J. A. (1988). *Signal Detection Theory and Psychophysics*. Los Altos: Peninsula Pub.

Hastie, R., and Kameda, T. (2005). The robust beauty of majority rules in group decisions. *Psychol. Rev.* 112, 494–508. doi:10.1037/0033-295X.112.2.494

Hong, L., and Page, S. E. (2008). Some microfoundations of collective wisdom. *Collect. Wisdom* 56–71.

Janis, I. L. (1972). *Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes*. Boston, MA: Houghton Mifflin Company.

Jose, V. R. R., Grushka-Cockayne, Y., and Lichtendahl, K. C. (2014). Trimmed opinion pools and the crowd's calibration problem. *Manage. Sci.* 60, 463–475. doi:10.1287/mnsc.2013.1781

Jose, V. R. R., and Winkler, R. L. (2008). Simple robust averages of forecasts: some empirical results. *Int. J. Forecast.* 24, 163–169. doi:10.1016/j.ijforecast.2007.06.001

Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Macmillan.

Kahneman, D., and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47, 263. doi:10.2307/1914185

Katsikopoulos, K. V., and King, A. J. (2010). Swarm intelligence in animal groups: when can a collective out-perform an expert? *PLoS ONE* 5:e15505. doi:10.1371/journal.pone.0015505

King, A. J., Cheng, L., Starke, S. D., and Myatt, J. P. (2012). Is the true 'wisdom of the crowd' to copy successful individuals? *Biol. Lett.* 8, 197–200. doi:10.1098/rsbl.2011.0795

Kording, K. P., and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature* 427, 244. doi:10.1038/nature02169

Körding, K. P., and Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* 10, 319–326. doi:10.1016/j.tics.2006.05.003

Krause, S., James, R., Faria, J. J., Ruxton, G. D., and Krause, J. (2011). Swarm intelligence in humans: diversity can trump ability. *Anim. Behav.* 81, 941–948. doi:10.1016/j.anbehav.2010.12.018

Kuczma, M., and Gilányi, A. (2009). *An Introduction to the Theory of Functional Equations and Inequalities: Cauchy's Equation and Jensen's Inequality*. Basel: Birkhäuser.

Kuncheva, L., Whitaker, C., Shipp, C., and Duin, R. (2003). Limits on the majority vote accuracy in classifier fusion. *Pattern Anal. Appl.* 6, 22–31. doi:10.1007/s10044-002-0173-7

Kurvers, R. H. J. M., Herzog, S. M., Hertwig, R., Krause, J., Carney, P. A., Bogart, A., et al. (2016). Boosting medical diagnostics by pooling independent judgments. *Proc. Natl. Acad. Sci. U.S.A* 113, 8777–8782. doi:10.1073/pnas.1601827113

Larrick, R. P., Mannes, A. E., and Soll, J. B. (2003). "The social psychology of the wisdom of crowds," in *Frontiers in Social Psychology: Social Judgment and Decision Making*, Vol. 6 (New York: Psychology Press), 227–242.

Larrick, R. P., and Soll, J. B. (2006). Intuitions about combining opinions: misappreciation of the averaging principle. *Manage. Sci.* 52, 111–127. doi:10.1287/mnsc.1050.0459

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi:10.1038/nature14539

Lima, S. L. (1995). Back to the basics of anti-predatory vigilance: the group-size effect. *Anim. Behav.* 49, 11–20. doi:10.1016/0003-3472(95)80149-9

List, C. (2011). The logical space of democracy. *Philos. Public Aff.* 39, 262–297. doi:10.1111/j.1088-4963.2011.01206.x

Liu, Y., and Yao, X. (1999). Ensemble learning via negative correlation. *Neural Netw.* 12, 1399–1404. doi:10.1016/S0893-6080(99)00073-8

Lorenz, J., Rauhut, H., Schweitzer, F., and Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. U.S.A* 108, 9020–9025. doi:10.1073/pnas.1008636108

Lorge, I., Fox, D., Davitz, J., and Brenner, M. (1958). A survey of studies contrasting the quality of group performance and individual performance, 1920–1957. *Psychol. Bull.* 55, 337–372. doi:10.1037/h0042344

Madirolas, G., and de Polavieja, G. G. (2015). Improving collective estimations using resistance to social influence. *PLoS Comput. Biol.* 11:e1004594. doi:10.1371/journal.pcbi.1004594

Mann, R. P., and Helbing, D. (2017). Optimal incentives for collective intelligence. *Proc. Natl. Acad. Sci. U.S.A* 114, 5077–5082. doi:10.1073/pnas.1618722114

Mannes, A. E. (2009). Are we wise about the wisdom of crowds? The use of group judgments in belief revision. *Manage. Sci.* 55, 1267–1279. doi:10.1287/mnsc.1090.1031

Mannes, A. E., Soll, J. B., and Larrick, R. P. (2014). The wisdom of select crowds. *J. Pers. Soc. Psychol.* 107, 276–299. doi:10.1037/a0036677

Manski, C. F. (2016). Interpreting point predictions: some logical issues. *Found. Trends Account.* 10, 238–261. doi:10.1561/1400000047

Marshall, J. A. (2011). "Optimal voting in groups with convergent interests," in *AAAI Spring Symposium: Modeling Complex Adaptive Systems as if They Were Voting Processes*, Menlo Park.

Marshall, J. A., Brown, G., and Radford, A. N. (2017). Individual confidence-weighting and group decision-making. *Trends Ecol. Evol.* 32, 636–645. doi:10.1016/j.tree.2017.06.004

Moussaïd, M., Perozo, N., Garnier, S., Helbing, D., and Theraulaz, G. (2010). The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE* 5:e10047. doi:10.1371/journal.pone.0010047

Myers, D. G., and Lamm, H. (1976). The group polarization phenomenon. *Psychol. Bull.* 83, 602. doi:10.1037/0033-2909.83.4.602

O'Neil, M., and Ryan, C. (2003). "Grammatical evolution," in *Grammatical Evolution* (New York: Springer), 33–47.

Page, S. E. (2008). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. New Jersey: Princeton University Press.

Polikar, R. (2012). "Ensemble learning," in *Ensemble Machine Learning*, eds C., Zhang and Y., Ma (Boston, MA: Springer US), 1–34.

Prelec, D., Seung, H. S., and McCoy, J. (2017). A solution to the single-question crowd wisdom problem. *Nature* 541, 532–535. doi:10.1038/nature21054

Raykar, V. C., Yu, S., Zhao, L. H., Hermosillo Valadez, G., Florin, C., Bogoni, L., et al. (2010). Learning from crowds. *J. Mach. Learn. Res.* 11, 1297–1322.

Reid, C. R., MacDonald, H., Mann, R. P., Marshall, J. A., Latty, T., and Garnier, S. (2016). Decision-making without a brain: how an amoeboid organism solves the two-armed bandit. *J. R. Soc. Interface* 13, 20160030. doi:10.1098/rsif.2016.0030

Roberts, G. (1996). Why individual vigilance declines as group size increases. *Anim. Behav.* 51, 1077–1086. doi:10.1006/anbe.1996.0109

Rokach, L. (2010). Ensemble-based classifiers. *Artif. Intell. Rev.* 33, 1–39. doi:10.1007/s10462-009-9124-7

Rosenthal, S. B., Twomey, C. R., Hartnett, A. T., Wu, H. S., and Couzin, I. D. (2015). Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion. *Proc. Natl. Acad. Sci. U.S.A* 112, 4690–4695. doi:10.1073/pnas.1420068112

Rothenberg, T. J., Fisher, F. M., and Tilanus, C. B. (1964). A note on estimation from a cauchy sample. *J. Am. Stat. Assoc.* 59, 460. doi:10.2307/2282999

Rudemo, M. (1982). Empirical choice of histograms and kernel density estimators. *Scand. J. Stat.* 1, 65–78.

Salganik, M. J., and Levy, K. E. C. (2015). Wiki surveys: open and quantifiable social data collection. *PLoS ONE* 10:e0123483. doi:10.1371/journal.pone.0123483

Schmidt, M., and Lipson, H. (2009). Distilling free-form natural laws from experimental data. *Science* 324, 81–85. doi:10.1126/science.1165893

Silver, N. (2012). *The Signal and the Noise: Why so Many Predictions Fail – but Some don't*. New York, NY: Penguin Press.

Simmons, J. P., Nelson, L. D., Galak, J., and Frederick, S. (2011). Intuitive biases in choice versus estimation: implications for the wisdom of crowds. *J. Consum. Res.* 38, 1–15. doi:10.1086/658070

Soll, J. B., and Larrick, R. P. (2009). Strategies for revising judgment: how (and how well) people use others' opinions. *J. Exp. Psychol. Learn. Mem. Cogn.* 35, 780–805. doi:10.1037/a0015145

Stein, B. E., and Stanford, T. R. (2008). Multisensory integration: current issues from the perspective of the single neuron. *Nat. Rev. Neurosci.* 9, 255–266. doi:10.1038/nrn2331

Sumpter, D. J. T. (2010). *Collective Animal Behavior*. New Jersey: Princeton University Press.

Sumpter, D. J. T., Krause, J., James, R., Couzin, I. D., and Ward, A. J. W. (2008). Report consensus decision making by fish. *Curr. Biol.* 18, 1773–1777. doi:10.1016/j.cub.2008.09.064

Sun, C., Shrivastava, A., Singh, S., and Gupta, A. (2017). Revisiting Unreasonable Effectiveness of Data in Deep Learning Era. *arXiv preprint arXiv:1707.02968*.

Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. New York: Doubleday.

Taleb, N. N. (2013). *Antifragile: Things That Gain from Disorder*. New York: Random House.

Tawn, J. A. (1988). Bivariate extreme value theory: models and estimation. *Biometrika* 75, 397–415. doi:10.1093/biomet/75.3.397

Treynor, J. L. (1987). Market efficiency and the Bean Jar experiment. *Financ. Anal. J.* 43, 50–53. doi:10.2469/faj.v43.n3.50

Trimmer, P. C., Houston, A. I., Marshall, J. A., Mendl, M. T., Paul, E. S., and McNamara, J. M. (2011). Decision-making under uncertainty: biases and Bayesians. *Anim. Cogn.* 14, 465–476. doi:10.1007/s10071-011-0387-4

Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131. doi:10.1126/science.185.4157.1124

Ungar, L., Mellors, B., Satopää, V., Baron, J., Tetlock, P., Ramos, J., et al. (2012). The good judgment project: a large scale test of different methods of combining expert predictions. *AAAI Fall Symposium: Machine Aggregation of Human Judgment*.

Valentini, G., Ferrante, E., and Dorigo, M. (2017). The best-of-n problem in robot swarms: formalization, state of the art, and novel perspectives. *Front. Robot. AI* 4. doi:10.3389/frobt.2017.00009

Vicsek, T., and Zafeiris, A. (2012). Collective motion. *Phys. Rep.* 517, 71–140. doi:10.1016/j.physrep.2012.03.004

Ward, A. J. W., Herbert-Read, J. E., Sumpter, D. J. T., and Krause, J. (2011). Fast and accurate decisions through collective vigilance in fish shoals. *Proc. Natl. Acad. Sci. U.S.A* 108, 2312–2315. doi:10.1073/pnas.1007102108

Watts, D. J. (2011). *Everything Is Obvious: Once You Know the Answer*. New York: Crown Business.

Whalen, A., and Yeung, S. (2015). "Using ground truths to improve wisdom of the crowd estimates," in *CogSci Conference*, Pasadena.

Wolf, M., Kurvers, R. H., Ward, A. J., Krause, S., and Krause, J. (2013). Accurate decisions in an uncertain world: collective cognition increases true positives while decreasing false positives. *Proc. Biol. Sci.* 280, 20122777. doi:10.1098/rspb.2012.2777

Yaniv, I. (2004). Receiving other people's advice: influence and benefit. *Organ. Behav. Hum. Decis. Process* 93, 1–13. doi:10.1016/j.obhdp.2003.08.002

Yaniv, I., and Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgments. *Organ. Behav. Hum. Decis. Process* 103, 104–120. doi:10.1016/j.obhdp.2006.05.006

Zhou, Z.-H., Jiang, Y., Yang, Y.-B., and Chen, S.-F. (2002). Lung cancer cell identification based on artificial neural network ensembles. *Artif. Intell. Med.* 24, 25–36. doi:10.1016/S0933-3657(01)00094-X

APPENDIX

1. Formal Definition of Key Mathematical Terms

1.1. Convexity

A convex function of a variable x which we denote as $f(x)$ with derivative $f(x)'$ is convex if

$$f(x) \geq f(y)'(x - y) + f(y), \tag{A1}$$

for all closed intervals $[x, y]$ (Kuczma and Gilányi, 2009).

1.2. Jensen's Inequality

Let $f(x)$ be a convex function of x and let $p(x)$ be a probability distribution of the values of x. Jensen's inequality states that (Kuczma and Gilányi, 2009)

$$f(E(x)) \leq E(f(x)). \tag{A2}$$

Here, E is the expected value operator and the expectation is with respect to probability distribution $p(x)$.

1.3. Bias and Variance

Let $p(x)$ be the probability distribution of some continuous variable x. Samples from the distribution $p(x)$ are used to estimate the value of some quantity with the true value y. Let us first deal with the case of simple averaging. If the mean of the distribution $p(x)$ is given as $\mu = \int x p(x) dx = E(x)$, then bias in given by $b = \mu - y$ and the variance is given by $\sigma^2 = \int (x - \mu)^2 p(x) dx$ (Geman et al., 2008).

More generally, if we use an estimator g(x) to estimate the value of y, then bias is given by the combination of equations $\mu = \int g(x)p(x)dx$ and $b = \mu - y$. The variance is given by $\sigma^2 = \int (g(x) - \mu)^2 p(x)dx$. For quadratic cost function, the expected squared error $\epsilon^2 = E((g(x) - y)^2)$ is given by the equation

$$\epsilon^2 = b^2 + \sigma^2. \tag{A3}$$

The above equation makes it clear that low errors occur only when both bias and variance are low.

1.4. Condorcet Voting

We consider a group of N individuals (N an odd number to avoid ties), where each individual votes independently and has probability p of producing the correct answer. The probability that the majority vote produces a correct answer is given as:

$$p(correct) = \sum_{m=(N+1)/2}^{N} \frac{N!}{m!(N - m)!} p^m (1 - p)^{N-m}. \tag{A4}$$

The formula can be understood as a weighted sum of binomial coefficients, where $\frac{N!}{m!(N-m)!}$ counts the total number of distinct ways to achieve m correct answers out N opinions and the term $p^m(1-p)^{N-m}$ calculates the probability of any individual occurrence with m correct answers. Condorcet theorem proves that if $p > 0.5$, then $p(correct)$ tends to 1 as N tends to infinity. Modern proofs of the claim usually rely on the central limit theorem (Sumpter, 2010), but with electronic computers it is also easy to just calculate the numerical value of each term in the sum

and analyses are no longer restricted to focusing on asymptotic behavior.

2. Derivation of the Many-Eyes Model

In the many-eyes model, a crowd of N individuals can only be wrong if all N people fail to detect an approaching predator. The probability of a single individual missing the predator (conditioned on the predator being present) is given by $(1 - t)$. The probability of all N individuals being wrong is $(1 - t)^N$. Since the predator attacks in only 50% of the cases, then the rate of errors is given by $0.5(1 - t)^N$. Therefore, the rate of correct decisions for the group is given by $1 - 0.5(1 - t)^N$.

Under the majority vote decision rule, the majority correctly detects a predator only if the number of people detecting a predator exceeds $(N + 1)/2$ (assuming odd N). The probability of a correct detection if predator is present is given by $y = \sum_{i=(N+1)/2}^{N} \frac{N!}{(N-i)!i!} t^i (1 - t)^{N-i}$. In the 50% of cases where no predator attacks, the crowd is always correct. Therefore, the overall rate of correct decisions in the case of applying majority vote to the panic scenario is $p_{correct} = 0.5 + 0.5y$ (black curve in **Figure 3**). If $t < 0.5$ then y tends to zero as N increases, which means that $p_{correct}$ tends to 0.5 at high N.

When we compare the many-eyes model with independent Condorcet voting, we need to ensure that it is the group decision mechanisms and correlation which cause the differences and that expected individual performance is the same for both models. The probability of an individual being correct in the panic model is $p = p(\text{attack})p(\text{correct|attack}) + p(\text{no attack})p(\text{correct|no attack}) = \frac{1}{2}t + \frac{1}{2}1 = \frac{1+t}{2}$, where the convention $p(\text{correct}|x)$ stands for the probability of making a correct decision under condition x. To calculate the green curve in **Figure 3**, we ensured that p had the same value for the red and the green curve and then applied Condorcet theorem to that p value.

Let us now relax the assumption of no false alarms. We define two probabilities, p_1 and p_2, which specify the probability of any given individual thinking that it detected a predator for the case of no predator present and for the case of a predator present, respectively. In the left panel of **Figure A1**, we illustrate how performance decays as we raise the value of p_1 given that the group continues to use the full vote procedure. As we can see, rising values of p_1 clearly diminish group accuracy, especially at large group size values. The rate at which performance deteriorates may be reduced if we adapt our decision threshold together with the value of p_1 and N. For any given scenario, the group now only declares a predator present if the number of animals detecting the predator is greater than some value T (the full vote corresponds to the case T = 0), with T a function of N and p_1.

If the threshold value was T, then the probability of a correct decision was given as $p(correct) = \frac{1}{2}sumP(T, N, p1) + \frac{1}{2}(1 - sumP(T, N, p2))$, where $sumP(T, N, p) = \sum_{m=0}^{T} \frac{N!}{m!(N-m)!} p^m (1 - p)^{N-m}$. For the adaptive threshold method, the value of T was calculated as $T = \left\lfloor \left(N \frac{ln \frac{1-p_1}{1-p_2}}{ln \frac{(1-p_1)p_2}{(1-p_2)p_1}} \right) \right\rfloor$, where N is again the group size and $\lfloor x \rfloor$ indicates the floor value of x. This expression for T was derived from the condition that $sumP(T, N, p_1) > sumP(T, N, p_2)$.

FIGURE A1 | The effect of false alarms. Left panel: The effect of varying false alarm probability p_1 on performance of the full vote method (p_2 was fixed at 0.3 for all curves on the panel). Right panel: performance can be restored if we allow the decision threshold T to vary as group size N changes. Black curve demonstrates how even with false alarm probability $p_1 = 0.05$, we can rescue performance if we choose the optimal T for each N. The yellow curve gives the performance of the group if all members were statistically independent (Condorcet) so as to show that small false alarm probabilities do not disrupt the ability to make gains as compared to the independent agent scenario.

3. Modeling the Influence of Distracting Alternatives

In this section, we map both our distracter method and Condorcet voting onto a diffusion process. We then show that the distracter scenario is a generalization of Condorcet theorem which also guarantees higher performance for all values of $N > 1$ and $K > 2$.

Let us first restate our assumptions. Our decision-makers are randomly sampled from an infinite crowd, where a fraction k of individuals are experts who always vote correctly no matter how many alternatives they face. The remaining fraction $1 - k$ of uninformed individuals each choose one option among the K alternatives at random. The uninformed individuals choose independently from each other. Out of the K alternatives, two are credible candidates for the correct option while $K - 2$ are distracting alternatives. Only the central decision-maker knows which alternatives are the irrelevant distracters, the voters remain ignorant of their existence.

During each round, we sample N individuals and have them vote. After the vote, we discard all the opinions that landed on the $K - 2$ distracting alternatives. Our final decision will be chosen according to which of the two alternatives that we considered realistic candidates for an answer received more votes. If both alternatives receive equal support, then we toss a fair coin to determine our final opinion.

First, it is clear that for the case of $K = 2$, the probability that a randomly chosen individual gives the correct answer is $p = k + \frac{1}{2}(1 - k)$. The voting under this scenario is equivalent to regular Condorcet voting since we have no irrelevant alternatives. This scenario will act as our baseline. We now show that as we increase K to values larger than 2, we outperform this baseline (for any value of $N > 1$).

We can view our voting procedure as a diffusion process. We are interested in the value δx, which measures the difference between the number of votes casted for the correct alternative relative to the number of votes casted for the incorrect alternative. After N votes have been cast, if $\delta x > 0$ then we have made the right choice. If $\delta x = 0$, we choose correctly with probability $\frac{1}{2}$. Otherwise we make a mistake.

During a single round of voting, δx acts as a random variable. With probability k, we sample an expert and δx increases by 1. In the $1 - k$ cases, where we miss the expert, we have two mutually exclusive alternatives. With probability $\frac{2}{K}(1 - k)$, we add to δx a random variable s which has value 1 with probability $\frac{1}{2}$ and value -1 also with probability $\frac{1}{2}$. This corresponds to a case where one of the uninformed individuals lands a vote among one of the two credible alternatives. With probability $\frac{K-2}{K}(1 - k)$, the value of δx remains the same as the vote of an uninformed individual lands on one of the distracters.

We can see that δx evolves as the sum of three mutually exclusive variables: the signal variable, the noise variable and the neutral variable. The probability of sampling the signal variable remains the same for all values of K. But the probability of sampling the noise variable decreases as a function of K. This gives the intuition why Condorcet scenario of $K = 2$ gives the worst performance. It happens because the signal-to-noise ratio is at its lowest value. We next give a more formal proof of our statements.

As can be seen from previous discussions, the probability of sampling an expert's opinions remains unchanged as K varies. In what follows next, the values of k and N will be fixed. Therefore, for all values of K, we can write

$$p_w(K) = \sum_{q_e=0}^{N} p(w|q_e, K)p(q_e), \qquad (A5)$$

where $p_w(K)$ is the overall probability of the group making an incorrect decision for a fixed value of K, $p(q_e)$ is the probability that a sample of N opinions will contain q_e expert opinions, and $p(w|q_e, K)$ gives the probability of making an incorrect decision

given that the sample of N opinions contained q_e experts. Note that the quantity $p(w|q_e, K)$ depends on K.

The next step in our proof is to show that $p(w|q_e, K = 2) \geq p(w|q_e, K > 2)$ for all values of q_e. Essentially, we will show that no matter how many experts a particular sample contained, adding more irrelevant alternatives always reduces or leaves the probability of error the same. The conclusion then follows immediately from considering equation (A2).

For $K = 2$ and q_e, we know that the number of noise opinions is fixed at $q_n = N - q_e$. Therefore, $p(w|q_e, K = 2) = p(w|q_e, q_n = N - q_e)$. For values of K larger than 2, the number of noise variables for any given sample containing q_e experts is not fixed, but varies between 0 and $N - q_e$ depending on how many of the variables were neutral variables because the opinions landed among the irrelevant distracters. Therefore, for $K > 2$, we may write

$$p(w|q_e, K) = \sum_{q_n=0}^{N-q_e} p(q_n|q_e, K)p(w|q_n, q_e). \quad (A6)$$

From equation (A3), we can see that the inequality $p(w|q_e, K = 2) \geq p(w|q_e, K > 2)$ holds as long as $p(w|q_n, q_e)$ is a non-decreasing function of q_n since $\sum_{q_n=0}^{N-q_e} p(q_n|q_e, K) = 1$.

The last step of our proof is to show that $p(w|q_n, q_e)$ is indeed a non-decreasing function of q_n. Let us compare the value of $p(w|q_n, q_e)$ with $p(w|q_n + 1, q_e)$. We can write $p(w|q_n, q_e)$ as

$$p(w|q_n, q_e) = \frac{1}{2}p(s_n = -q_e|q_n) + \sum_{s_n=-q_n}^{-q_e-1} p(s_n|q_n), \quad (A7)$$

where s_n is a random variable that is calculated as the sum of q_n randomly and independently sampled noise variables which each take the values 1, -1 with probability 0.5. This equation is a simple application of the idea that in order to overturn the correct signal induced by the q_e experts, the q_n noise variables must have a sum equal to or lower than the value $-q_e$. The term $p(s_n = -q_e|q_n)$ contributes half its value because ties are broken by a coin toss.

We can relate $p(w|q_n, q_e)$ to $p(w|q_n + 1, q_e)$ by noting that in any random sample, when moving from q_n to $q_n + 1$, we are simply adding a number -1 or a number $+1$ to the value of s_n already present in the sum. If the value of s_n was already lower than $-q_e - 2$, then the addition of even a $+1$ is not enough to overwhelm the destructive influence of the noise. Also, if s_n was already higher than $-q_e + 2$ then even sampling a -1 is not enough to overturn the signal. Therefore, the only terms that may have any effect concern the boundary cases of $s_n = -q_e + 1$, $s_n = -q_e$, and $s_n = -q_e - 1$. Putting all this together,

$$
\begin{aligned}
p(w|q_n + 1, q_e) = {} & \frac{1}{4}p(s_n = -q_e + 1|q_n) + \frac{1}{2}p(s_n = -q_e|q_n) \\
& + \frac{3}{4}p(s_n = -q_e - 1|q_n) \\
& + \sum_{s_n=-q_n}^{-q_e-2} p(s_n|q_n) = p(w|q_n, q_e) \\
& + \frac{1}{4}p(s_n = -q_e + 1|q_n) \\
& - \frac{1}{4}p(s_n = -q_e - 1|q_n). \quad (A8)
\end{aligned}
$$

We are left to examine the term $\frac{1}{4}p(s_n = -q_e + 1|q_n) - \frac{1}{4}p(s_n = -q_e - 1|q_n)$, which turn out to be non-negative for all values of q_e, q_n. If $q_e = 0$, then the term is obviously zero because the distribution of s_n is symmetric and $p(w|q_n, q_e) = p(w|q_n + 1, q_e)$. A similar conclusion holds if $q_e > q_n$ because then both probabilities of the difference term are zero. The more interesting case concerns $0 < q_e < q_n + 1$. In that case $\frac{1}{4}p(s_n = -q_e + 1|q_n) > \frac{1}{4}p(s_n = -q_e - 1|q_n)$, because the distribution of s_n peaks at zero and decreases monotonically as s_n decreases away from 0. The combination of the three cases gives us the proof that $p(w|q_n, q_e)$ is non-decreasing in q_n which concludes our proof.

For the limit of K tends to infinity, we can give a surprising and compact expression for how performance varies with group size. For large K, nearly all uninformed opinions land on the distracter alternatives. Therefore, the only way a mistake will occur is if no experts happen to be selected into the group and the coin flip favors the wrong alternative. The probability of such an event is $\frac{1}{2}(1 - k)^N$ and, therefore, the probability of getting the correct answer is $1 - \frac{1}{2}(1 - k)^N$ which is the same equation as we had for the many-eyes model but with k replaced by t.

4. Training Neural Networks

The neural networks were trained in Tensorflow (Abadi et al., 2016). For the first task, we used an input layer of size 30 and two hidden layers of size 75 with rectified linear activation. The cost function optimized was the mean square error. We created a simulated dataset as described in main text with 5,000 training examples. The network was trained with ADAM using learning rate of 0.01.

For the case of doctor's estimations of presence of skin cancer, we used a dataset comprised of evaluations of 40 doctors on 108 different cases of potential melanomas from Kurvers et al. (2016). We split these cases into 54 training, 27 validation, and 27 test cases. For each groups size, we trained 50 different networks using 50 different random partitions of the data into 54, 27, and 27 cases.

Accuracy of each doctor was determined computing her Youden's index (J = sensitivity + specificity − 1) over the training cases. We then produced all 780 combinations of 2 doctors, and 1,000 random combinations of 3, 5, and 7 doctors, and computed the accuracy of each group using the different heuristics proposed over the test cases. The performance of the heuristics for each group size was then determined by averaging its value across all groups.

To train each network, we generated 54 training instances combining judgments, accuracies, and confidence ratings of each random group on each particular case. For example, for groups of 2 doctors each input was then composed of accuracy of first doctor, confidence of first doctor, accuracy of second doctor, and confidence of second doctor. Accuracies were multiplied by −1 if the doctor had judged the case as negative. We used the training cases to train the network and the validation cases to select the state of the network that produced the best performance. Then this particular state was applied to make predictions over the test cases and to compare its performance with the heuristics applied to the pairs.

The network architecture was different for each group size. For groups of 2 and 3 doctors, two hidden layers were used; and for

groups of 5 and 7, only one hidden layer was used. The size of each layer was 250 for groups of 2, 5, and 7 doctors, and 100 for groups of 3 doctors. All hidden layers had rectified linear activation. The network was trained with ADAM using learning rate of 0.0001 for groups of 2 doctors and 0.00001 for groups of 3, 5, and 7.

The cost function was selected to match the accuracy measured by the Youden index. This index is of the form $J = \text{TP}/(\text{TP} + \text{FN}) + \text{TN}/(\text{TN} + \text{FP}) - 1$, with TP standing for true positives, FN for false negatives, TN for true negatives, and FP for false positives. As the output of the network was the probabilities p and $1 - p$ that the case fed was a positive or a negative, the expected value of the Youden's index would be of the form $E[J] = \sum_{i=1}^{n_p} \frac{p_i}{n_p} + \sum_{i=n_p+1}^{n_p+n_n} \frac{1-p_i}{n_n} - 1$, where n_p (n_n) is the number of positives (negatives) and the first (second) sum is over the positive (negative) cases. The cost function optimized was then of the form $0.5(1 - E[J])$, which is 0 at the maximum expected Youden's index ($E[J] = 1$) and 1 at the minimum Youden's index ($E[J] = -1$).

How to: Using Mode Analysis to Quantify, Analyze, and Interpret the Mechanisms of High-Density Collective Motion

Arianna Bottinelli [1]* and Jesse L. Silverberg [2, 3]*

[1] NORDITA, Stockholm University, Stockholm, Sweden, [2] Wyss Institute for Biologically Inspired Engineering, Harvard University, Boston, MA, United States, [3] Department of Systems Biology, Harvard Medical School, Boston, MA, United States

While methods from statistical mechanics were some of the earliest analytical tools used to understand collective motion, the field has substantially expanded in scope beyond phase transitions and fluctuating order parameters. In part, this expansion is driven by the increasing variety of systems being studied, which in turn, has increased the need for innovative approaches to quantify, analyze, and interpret a growing zoology of collective behaviors. For example, concepts from material science become particularly relevant when considering the collective motion that emerges at high densities. Here, we describe methods originally developed to study inert jammed granular materials that have been borrowed and adapted to study dense aggregates of active particles. This analysis is particularly useful because it projects difficult-to-analyze patterns of collective motion onto an easier-to-interpret set of eigenmodes. Carefully viewed in the context of non-equilibrium systems, mode analysis identifies hidden long-range motions and localized particle rearrangements based solely on the knowledge of particle trajectories. In this work, we take a "how to" approach and outline essential steps, diagnostics, and know-how used to apply this analysis to study densely-packed active systems.

Keywords: mode analysis, active matter, jammed active matter, collective motion, human crowds, soft spots, rattler, topological defects

1. INTRODUCTION

The vast complexity of human neurobiology gives rise to a rich interior life filled with thoughts, moods, motivations, ideas, discourse, and imagination. Given this lived experience, it's remarkable that the challenges for explaining an individual's specific actions recede when we instead consider emergent group-scale human collective behavior [1]. This observation has fueled a surge of interest at the intersection of social psychology, behavioral economics, and data science, resulting in highly-effective and systematic strategies for broad-based social engineering [2–4]. These behavioral interventions, often called "nudges," are a modern staple for organizations ranging from governments to Fortune 500 companies seeking to broadly reshape the individual decisions that give rise to emergent collective behavior [5, 6]. While nudges are straight-forward to implement when the collective behavior occurs frequently, low-probability high-impact "black swan" events [7, 8], such as disasters at mass gatherings, call for alternative strategies. For example, music concerts, religious pilgrimages, sporting competitions, political protests, and consumer shopping holidays

occasionally lead to spontaneous and shockingly injurious large-scale human collective motion [9–11]. In these situations, high crowd densities and limited communication can result in fatalities through stampedes, crowd crush, or escape panic. These negative outcomes offer substantial impetus for the development of preventative safety strategies and life-saving interventions that can be deployed at mass gatherings. With this goal in mind, we describe a physical and mathematical approach to understand, predict, and ultimately prevent tragic human collective motion. By unraveling the basic physical mechanisms of emergent collective motion in this complex system, we aim to ground and inspire future intervention strategies.

The methods for analyzing high-density human crowds described here stem from an uncanny resemblance between mass gatherings and disordered granular packings (**Figure 1**) [12–15]. In both cases, we observe dense, irregular structure that persists over extended periods of time, punctuated with large sudden collective motion or spatially localized rearrangements. The existing research on these complex materials provides a systematic framework for characterization of collective motion along with theoretical tools that connect local structure to dynamical response [16–22]. The method derives from an analysis of disordered linear systems at equilibrium wherein eigenvalues and eigenmodes of the local displacement correlation matrix convey information about structural stability [19]. In this framework, eigenmodes relate to the magnitude and directions of collective motion, while the eigenvalues correspond to the excitation energies. Displacements can then be expressed as a linear combination of these modes with time-dependent coefficients. To the extent that such approximations effectively describe non-equilibrium jammed active matter [20], we use this framework to study aggregated crowds and their penchant for collective motion. In the context of human gatherings, our results enable an understanding of specific mechanisms for dangerous collective motion and the physical mechanisms underlying crowd disasters [23].

In the sections that follow, we describe how to implement the basic steps of eigenmode analysis and effectively interpret the results for high-density human crowds. While the method itself is quite general, we demonstrate the protocol through the example of an asocial model for high-density human collective behavior. In step-by-step fashion, we specifically emphasize practical tips for working with data, whether it be numerically simulated or empirically collected.

2. METHODS

We divide the methods section into two subsections. The first subsection describes a physical model for asocial human collective behavior in high-density crowds. This model is motivated by previous work and provides a specific means for simulating individual human trajectories [23, 24]. The second subsection details an analysis protocol that takes trajectory data as input, and converts this information into a quantitative description of emergent collective behavior. This protocol may
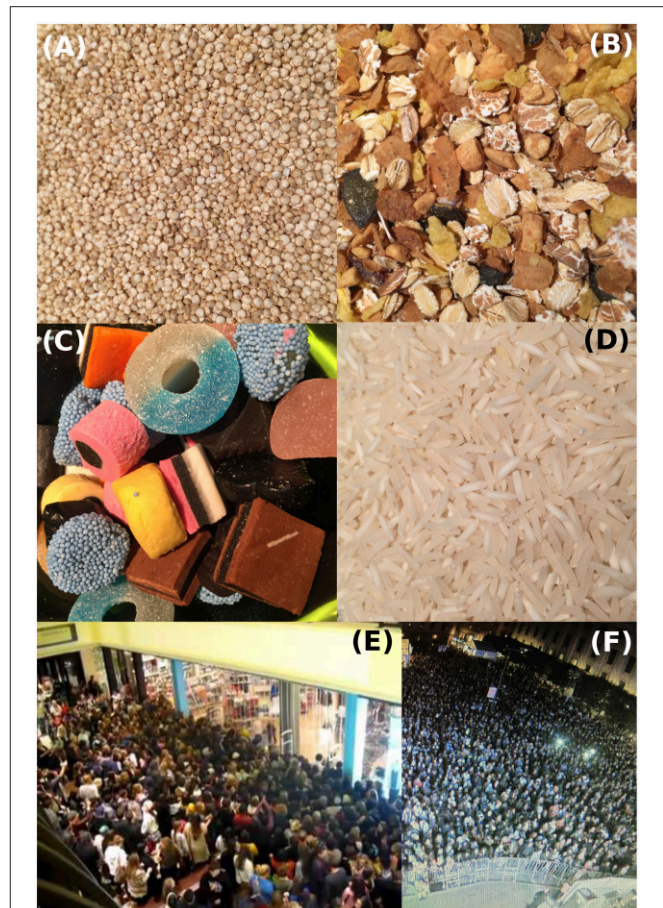
FIGURE 1 | Common examples of non-human granular media including **(A)** quinoa, **(B)** muesli, **(C)** mixed candy, and **(D)** rice, compared to high-density human crowds at **(E)** a Black Friday sales event and **(F)** an outdoor concert. Similarities with the former inspires an analysis of the latter. This collection of images represents a cross-section of the various geometries, heterogeneities, and interactions that fall within the purview of granular media methods.

be broadly applied to trajectory data from sources other than the asocial model.

2.1. A Physical Model to Study High-Density Human Collective Behavior

To quantitatively model human collective behavior, we take a Newtonian force-based approach to generate complex emergent phenomena [25–29]. While systems studied within this "active matter" framework are generally non-equilibrium, the resulting phenomenology is often reminiscent of behaviors analyzed in the fields of statistical mechanics, granular materials, and fluid dynamics. As such, there is a rich tradition of concepts from these fields intermixing [13–15].

2.1.1. Equations of Asocial Model

We specifically investigate human collective motion in high-density crowds. We therefore simplify the richness of human life to four forces, numerically simulate the resulting equations

of motion, and investigate the emergent collective behavior. Referring to these simplified humans as Self-Propelled Particles (SPPs), we assume they are all interested in going toward a single point of interest, $\mathcal{P}$. This point could be the front of a concert stage, the police line at a protest, or the exit of a stadium. While these activities have clear social differences, their commonality is in the accumulation of a large densely-packed crowd drawn by a common attraction. Each SPP indexed by i can now be described as a disk with radius r_0 positioned at a point $\vec{r}_i(t)$ and subject to: a pairwise soft-body repulsive collision force

$$\vec{F}_i^{\text{repulsion}} = \epsilon \sum_{j \neq i}^{N} \left(1 - r_{i,j}/2r_0\right)^{3/2} \hat{r}_{i,j}, \tag{1}$$

which is non-zero only when the distance between two particles $|\vec{r}_i - \vec{r}_j| = |d_{ij}\hat{r}_{i,j}| = d_{ij} < 2r_0$ [23, 24]; a self-propulsion force

$$\vec{F}_i^{\text{propulsion}} = \mu(v_0 \hat{p}_i - \vec{v}_i), \tag{2}$$

where v_0 is a constant preferred speed, $\vec{v}_i$ is the current velocity of the i^{th} SPP, and $\hat{p}_i$ is a unit vector pointing from each particle's center to the common point of interest $\mathcal{P}$; a randomly fluctuating force with components

$$\vec{F}_i^{\text{noise}} = \vec{\eta}_i, \tag{3}$$

drawn from a zero-mean Gaussian distribution and standard deviation σ defined by the correlation function $\langle \eta_{i,\lambda}(t)\eta_{i,\kappa}(t')\rangle = 2\mu^{-1}\sigma^2\delta_{\lambda\kappa}\delta(t-t')$, ensuring noise is spatially and temporally decorrelated; and finally, a wall collision force used to construct a confining simulation environment

$$\vec{F}_i^{\text{wall}} = \epsilon \sum_{\text{walls}} \left(1 - r_{i,w}/r_0\right)^{3/2} \hat{w}, \tag{4}$$

which is pointed along each wall's outward normal direction $\hat{w}$ and non-zero when the distance of the SPP from the wall $r_{i,w} < r_0$. While other repulsion forces have been used in similar models of human collective behavior [30–32], the functional form of $\vec{F}_i^{\text{repulsion}}$ and $\vec{F}_i^{\text{wall}}$ used here comes from treating SPP collisions as a Hertzian contact mechanics problem involving frictionless elastic spherical bodies [24, 33, 34]. Summing the forces from Equations (1) to (4), we find the evolution of each SPPs dynamics is driven by

$$\ddot{\vec{r}}_i(t) = \vec{F}_i^{\text{repulsion}} + \vec{F}_i^{\text{propulsion}} + \vec{F}_i^{\text{noise}} + \vec{F}_i^{\text{wall}} = \vec{F}_i^{\text{total}}, \tag{5}$$

where the relative magnitudes of individual force terms can be tuned through the scalar coefficients ϵ and μ. Because Equation (5) lacks any terms that reflect social interaction, we refer to it as an *asocial model* for high-density human collective behavior [23].

2.1.2. Numerical Implementation
Simulations take place in a rectangular room with wall length $L \gg r_0$ centered at the origin $(0, 0)$. The attraction point $\mathcal{P}$ is

placed at the center of the right wall at $(L/2, 0)$ and N SPPs are seeded at random initial positions with zero initial speed and acceleration. At every integration time-step we compute the total force acting on individual SPPs at their current positions $\vec{r}_i(t)$ and evolve their trajectories according to Equation (5) (**Figure 2**). This calculation is performed numerically with the Newton-Stomer-Verlet algorithm, which finds the *next* position using the *current* velocity $\dot{\vec{r}}_i(t) = \vec{v}_i(t)$ and *current* acceleration $\ddot{\vec{r}}_i(t) = \vec{F}_i^{\text{total}}(t)$ according to:

$$\vec{r}_i(t + \Delta T) = \vec{r}_i(t) + \vec{v}_i(t)\Delta T + \frac{1}{2}\vec{F}_i^{\text{total}}(t)(\Delta T)^2. \tag{6}$$

The *next* acceleration $\ddot{\vec{r}}_i(t + \Delta T) = \vec{F}_i^{\text{total}}(t + \Delta T)$ is then calculated at this new position, so that the *next* velocity can be determined by an average of the *current* and *next* accelerations:

$$\vec{v}_i(t + \Delta T) = \vec{v}_i(t) + \frac{\vec{F}_i^{\text{total}}(t) + \vec{F}_i^{\text{total}}(t + \Delta T)}{2}\Delta T. \tag{7}$$

Looping this sequence of calculations produces trajectories for each of the N SPPs (**Figure 2**). In Equations (6) and (7) the parameter Δt defines the integration step, which should be small enough to ensure smooth trajectories, but large enough to achieve reasonable computation times. Here, we run our simulations for $t = 3{,}000\tau$ units of time with $\Delta T = \tau/10 = 0.10$, yielding a total of 30,000 integration steps. Every 10 time steps we record each SPPs position $\vec{r}_i(t)$ as well as the pressure due to radial contact forces

$$P_i(t) = \frac{1}{2\pi r_0}\left[\vec{F}_i^{\text{repulsion}} + \vec{F}_i^{\text{wall}}\right] \cdot \hat{r}_i, \tag{8}$$

where the dot product is with the unit normal vector centered on the i^{th} SPP. We consistently find transient motion dominates the first $\approx 50\tau$ of the simulation resulting in far-from-equilibrium effects (**Figures 2**, **3**, linear path segments). By 300τ the SPPs have aggregated near $\mathcal{P}$ and formed a stable, dense, disordered aggregate with $\vec{F}_i^{\text{propulsion}}$ acting similar to an external field confining the SPPs. Within the aggregate, collision and noise forces are responsible for position fluctuations, causing each particle to move randomly around its average position (**Figure 3** inset, densely accumulated trajectory data near the point of interest $\mathcal{P}$ denoted by $\star$).

2.1.3. Model Parameters and Time Scales
Setting model parameters in terms of the fundamental simulation unit length ℓ and unit time τ allows us to maintain careful control over the relative force and time scales while not explicitly committing to dimensionful units such as meters or seconds. As such, we simulate $N = 200$ SPPs of radius $r_0 = \ell/2$ in a region of size $L = 50\ell$. These choices ensure the simulation box size L is larger than the typical aggregated SPP group size $\sim \ell\sqrt{N}$. We also set the SPP preferred speed $v_0 = \ell/\tau$, the random force standard deviation $\sigma = \ell/\tau^2$, and the force scale coefficients $\epsilon = 25\ell/\tau^2$, $\mu = \tau^{-1}$ [24]. For our analysis, we run 10 independent simulations of the dynamics with this set of parameters and random initial conditions.

FIGURE 2 | Four screenshots of a typical simulation run for our asocial model of high-density human crowds. Here, we see $N = 200$ SPPs aggregating near a point of interest $\mathcal{P}$ denoted by $\star$, which is located at the right-most edge of a simulation box. SPPs self-organize into a dense and disordered aggregate, where shading is assigned according to the radial pressure at time t. The rectangular outline for $t = 15$ gives an impressionistic sense of the simulation box; the true border is taller and somewhat wider.

FIGURE 3 | Trajectories $r_i(t)$ for $N = 200$ SPPs in a typical simulation run of the asocial model for high-density human crowds. After a transient period (radial lines toward $\star$), the SPPs reach a stable state where they randomly oscillate about their average position (inset, dense squiggles). Each SPP is given a different color to more easily distinguish it from its neighbors.

Setting $r_0 = \ell/2$ means that in the absence of any other interactions a SPP would move a distance equal to its diameter in the time τ. This choice approximates relaxed pedestrian motion if we were to have τ equal to one second [30].

The coefficient μ relates to an exponential relaxation time for the self-propulsion force, which can be seen by solving for free acceleration of a SPP. Specifically, $dv/dt = \mu(v_0 - v)$ has a solution $v(t) = v_0[1 - \exp(-\mu t)]$ when $v(t = 0) = 0$. This expression shows that an unobstructed SPP will exponentially approach its preferred speed v_0 with a timescale μ^{-1}.

Both SPP-SPP and SPP-wall collisions are subject to a Hertzian repulsion force directed normally to the surface of contact with a magnitude set by the coefficient ϵ. These forces are non-singular, making them numerically stable and easy to simulate, which is particularly useful for studies of jammed soft granular matter [34]. In the context of human crowds, Equations (1) and (4) greatly simplify collisional interactions, while still capturing the fact that people can be partially compressed with a rising non-linearity as the stresses increase. Here, ϵ is a constant for all SPPs equal to $25\ell/\tau^2$, which guarantees the collision force scale is substantially larger than the self-propulsion force scale $\mu v_0 = \ell/\tau^2$.

To ensure collisions and random force fluctuations contribute roughly equally to the SPPs dynamics, the collision time scale τ_{coll} and the random force time scale τ_{noise} must be comparable. In our asocial model, the random collision time scale $\tau_{\text{coll}} = 1/(2r_0v_0n) \approx (\pi/4)\tau$ is given by the mean-free path $(2r_0n)^{-1} \approx (\pi/2)r_0$ divided by the preferred speed v_0. The average crowd density $n \approx N/\pi(\sqrt{N}r_0)^2$ is estimated by noting the steady-state configuration of SPPs is roughly a half-circle with radius

$\sqrt{N}r_0$ surrounding $\mathcal{P}$. The noise time scale τ_{noise} can be found by calculating the amount of time required for random fluctuations to change the correlation function $\langle [v_i(\tau_{\text{noise}}) - v_i(0)]^2 \rangle = 2\mu^{-1}\sigma^2\tau_{\text{noise}}$ by an amount equal to v_0^2. Hence, $\tau_{\text{noise}} = \mu v_0^2/2\sigma^2$. Because the unit speed v_0 is fixed by the fundamental simulation units, and μ is set by the self-propulsion relaxation time, we simply let $\sigma = \ell/\tau^2$ to satisfy $\tau_{\text{noise}} = \tau/2 \approx \tau_{\text{coll}}$ at steady-state.

2.2. Analysis Protocol of Trajectory Data

In section 2.1, we outlined a physical model for generating trajectory data of simulated human crowds using SPPs. In this section, we provide a step-by-step protocol for analyzing trajectory data using mode analysis as a means for predicting emergent collective motion. While we demonstrate the protocol with simulated data from the asocial model, our analysis can be applied equally well to other active jammed systems. As such, we cast our discussion in terms of "agents," which could be either SPPs, actual humans being studied with cameras, or other examples of high-density active matter under consideration.

2.2.1. Step 1: Calculate the Displacement Covariance Matrix Components C_{ij} to Estimate the Ground-Truth Correlation Matrix C_p

Each of the trajectories $\vec{r}_i(t) = \langle x_i(t), y_i(t) \rangle$ provide spatially resolved position data on the $i = 1, \ldots, N$ agents at discrete time points t. From these trajectories, we want to determine the displacement correlation matrix C_p, which contains information about pair-wise correlated motion arising from the local interactions and the resulting collective motion. Ideally, this matrix is a statistical quantity averaged over all realizations of the underlying system. In practice, there are a limited number of computational runs or experimental measurements available. Thus, we calculate the covariance matrix $[C_{ij}]$, whose components C_{ij} converge to the ground-truth correlation matrix C_p in the $t \to \infty$ limit. These components are

$$C_{ij} = \left\langle \left[\vec{r}_i(t) - \langle \vec{r}_i \rangle \right] \cdot \left[\vec{r}_j(t) - \langle \vec{r}_j \rangle \right] \right\rangle, \quad (9)$$

where time averages $\langle \cdots \rangle$ are calculated for each realization of the underlying random process over a statistically-independent sub-sampling of the time series data. Conventionally, this is the equivalence between time and ensemble averaging. *Critical Note 1:* While Equation (9) is standard notation in the literature, it obscures the fact that there are actually two sets of covariance matrix components to calculate: one each for the x and y directions. A more transparent representation would be

$$C_{ij}^x = \left\langle [x_i(t) - \langle x_i \rangle] \cdot [x_j(t) - \langle x_j \rangle] \right\rangle \quad \text{and}$$

$$C_{ij}^y = \left\langle [y_i(t) - \langle y_i \rangle] \cdot [y_j(t) - \langle y_j \rangle] \right\rangle. \quad (10)$$

However, the cumbersome nature of these expressions tends to favor the aesthetics and compactness of Equation (9). *Critical Note 2:* Time-averaging calls for a judicial eye that balances two competing demands: sub-sampling in time should be spaced out to reduce effects from auto-correlated motion,

while simultaneously leaving a sufficient number of statistically-independent "snap-shots" of the system for the components of the covariance matrix $[C_{ij}]$ to converge to the correlation matrix C_p. A straight-forward convergence criteria is to ensure the ratio $2N/N_t < 1.5$, where N_t is the number of independent snap-shots [16]. *Critical Note 3:* Because mode analysis depends on the eigenvalues and eigenmodes of $[C_{ij}]$, the quality of time averaging can be tested by examining how the eigenvalue spectrum converges as a function of the temporal sampling rate [16]. Specifically, the spectrum should be insensitive to the sampling rate; comparing spectra generated at different sampling rates will indicate whether a specific value is in this regime. *Critical Note 4:* Many physical models can be treated with scaling analysis to determine a minimal estimate for the time scale of auto-correlated motion. In the asocial model described in section 2.1, noise and repulsive collision forces dissipate auto-correlation. Thus, the trajectory data should be sub-sampled at temporal intervals spaced out longer than τ_{noise} and τ_{coll}. *Critical Note 5:* In some circumstances such as an external perturbation or sudden motion, the agents being studied can undergo rearrangements in position that change the internal structure of which agents are in contact with each other. When this occurs, the covariance matrix components C_{ij} must be recalculated from post-rearrangement trajectory data and convergence of $[C_{ij}] \to C_p$ must be rechecked.

Example: Determining appropriate temporal sub-sampling of asocial model. For the parameter values described above, we find the simulation reaches steady state after $\approx 50\tau$. We generously discard the first 300τ to eliminate far-from-equilibrium transients leaving $2,700\tau$ of trajectory data for each SPP. Because $\tau_{\text{noise}} = \tau/2 \approx \tau_{\text{coll}}$, then $\tau/2$ is the minimal estimate for temporal sub-sampling. From this baseline, we check convergence of the eigenvalue spectrum for intervals longer than $\tau/2$, and find a temporal sub-sampling of 10τ is adequate. We are then left with $N_t = (2,700\tau)/(10\tau) = 270$ statistically-independent snap-shots of the system to use when calculating temporal averages in Equation (9), which is consistent with the $2N/N_t = 1.48 < 1.5$ criteria.

2.2.2. Step 2: Calculate the Eigenmodes and Eigenvalue Spectrum

Having calculated a displacement covariance matrix $[C_{ij}]$ that approximates the ground-truth correlation matrix C_p, we can now use standard numerical techniques to compute the eigenvalues λ_m and their corresponding eigenmodes $\vec{e}_m$. The index $m = 1, \ldots, N$ runs over all agents under consideration. Often, the terminology "eigenmode" is simply shortened to "mode," which is a callback to the field's roots in analyzing harmonic vibrational motion. *Critical Note 1:* Sorting the eigenvalues in decreasing order and plotting as a function of index m gives the *spectrum* of the covariance matrix, which in the $t \to \infty$ limit converges to the spectrum of the correlation matrix. *Critical Note 2:* As in the previous step, notational conventions obscure the fact that there are two sets of eigenvalues and eigenmodes. Specifically, both x and y directions have their own eigenvalues λ_m^x and λ_m^y for a total of $2N$ eigenvalues

with two distinct spectra. Likewise, each eigenmode is a two-dimensional displacement vector field. For example, the $m = 1$ eigenmode is more transparently expressed as $\langle \vec{e}_1^x, \vec{e}_1^y \rangle$, with $\vec{e}_1^x = \{e_{1,1}^x, e_{1,2}^x, \ldots, e_{1,N}^x\} = \{e_{1,j}^x\}$ and $\vec{e}_1^y = \{e_{1,1}^y, e_{1,2}^y, \ldots, e_{1,N}^y\} = \{e_{1,j}^y\}$. These $2N$ components indexed by $j = 1, \ldots, N$ are the $m = 1$ eigenmode values in the x and y directions for each of the j agents. *Critical Note 3:* The eigenmodes are typically normalized such that $\sum_{j=1}^{N} e_{m,j}^x = 1$, $\sum_{j=1}^{N} e_{m,j}^y = 1$, and the norm $\sum_{j=1}^{N} |\vec{e}_{m,j}|^2 = 2$. *Critical Note 4:* Within a linear approximation, the eigenmodes express relative magnitude and directions for harmonic oscillation of the agents about their mean positions. The eigenvalues relate to the corresponding excitation energies of these modes. Trajectory data $\vec{r}(t)$ can then be decomposed into a linear combination of these modes with a collection of $2N$ time-dependent coefficients. As the non-linearity and non-equilibrium nature of the system asserts itself, these approximations breakdown.

2.2.3. Step 3: Find and Remove Rattlers

Plotting low-m eigenmodes as two-dimensional displacement fields frequently reveals a small number of agents $N_r \ll N$ that represent a disproportionately large amount of the overall motion. This phenomenon is well-known to arise in both experimental and simulated jammed systems, where such agents are called "rattlers." While most often studied in colloidal suspensions and vibrated/sheared granular packings, rattlers get their name because they are surrounded by highly-constrained neighbors that create a rigid cage enclosing enough space for the rattler to freely move about [22, 35]. This under-constrained motion is therefore a consequence of local structure. Because (i) rattlers tend not to participate in global collective motion and (ii) the location of rattlers is impossible to predict *a priori*, we must identify them *a posteriori*, remove the N_r rattlers from consideration, and recompute Equation (9) with the subset of $N - N_r$ agents. A threshold identification criteria is useful for systematically finding rattlers. In particular, for each mode m we check for agents with index j whose individual displacement

$$|\vec{e}_{m,j}| = \left[(e_{m,j}^x)^2 + (e_{m,j}^y)^2 \right]^{1/2} \tag{11}$$

is greater than the mode's average displacement

$$\langle |\vec{e}_m| \rangle = \frac{1}{N} \sum_{j=1}^{N} \left[(e_{m,j}^x)^2 + (e_{m,j}^y)^2 \right]^{1/2} \tag{12}$$

plus ξ_r times the mode's standard deviation

$$\sigma_m = \left[(\sigma_m^x)^2 + (\sigma_m^y)^2 \right]^{1/2}$$
$$= \left[\sum_{j=1}^{N} \frac{[e_{m,j}^x - \langle |e_m^x| \rangle]^2 + [e_{m,j}^y - \langle |e_m^y| \rangle]^2}{N - 1} \right]^{1/2}. \tag{13}$$

More succinctly, rattlers are defined as agents with index j with $|\vec{e}_{m,j}| \geq \langle |\vec{e}_m| \rangle + \xi_r \sigma_m$, for a given threshold value ξ_r.

Once identified, these N_r agents can be removed from further consideration. *Critical Note 1:* The removal of rattlers ensures low-m modes reflect genuine collective motion of the system instead of the free vibrations of under-constrained agents. *Critical Note 2:* Selection of an appropriate threshold should be done by testing multiple values of ξ_r and examining how the fraction N_r/N varies. If the threshold is too high, then the eigenmodes will continue to be dominated by the uncorrelated motion of rattlers. If the value is too low, then the analysis risks under-sampling the calculation of C_{ij} and no collective motion will be detected. As a general heuristic, the value of N_r arising from ξ_r should be the minimum value that removes anomalous uncorrelated motion from the low-m modes. No single value of N_r/N will be appropriate for all circumstances, and the selection of ξ_r should be given due consideration. In the next step we provide further useful information to aid in the choice of ξ_r and illustrate the procedure in the case of our simulation data.

2.2.4. Step 4: Re-calculate the Eigenmodes and Eigenvalue Spectrum without Rattlers

Having provisionally identified the N_r rattlers in step 3, we repeat steps 1 and 2 to re-calculate the eigenmodes and eigenvalue spectrum from $[C_{ij}]$ with the remaining $N - N_r$ agents. When examining the new modes and spectra produced by step 2, we generally find the first attempt at removing rattlers is insufficient and steps 1 through 3 should be performed as an iterative process to optimally determine ξ_r. *Critical Note 1:* Concretely, this strategy starts with an initial guess for the threshold ξ_r, and determines a final value by qualitatively and quantitatively checking how ξ_r affects the eigenmodes and eigenvalue spectrum. The goal is to find a value of ξ_r that filters rattlers from the overall collective motion in low-m eigenmodes. *Critical Note 2:* While iterating, it's useful to: (i) visually confirm whether the eigenmode plots are being affected by under-constrained motions, and (ii) confirm the eigenvalue spectra retain their basic shape. This information is essential for making an informed decision on the next threshold value to test. *Critical Note 3:* A reasonable initial guess for the rattler threshold is to set $\xi_r = 1$ with the expectation that the final value will be larger.

Example: Iteratively selecting ξ_r. Working with trajectory data generated by the asocial model, we perform steps 1 through 4 of the analysis protocol and test various values for the rattler threshold ξ_r (**Figure 4**). In this instance, we seek to find and remove rattlers from the first 10 modes by examining values of ξ_r ranging from 1 to 6. As prescribed in the protocol, we calculate and re-calculate the eigenvalue spectrum eliminating provisionally identified rattlers from consideration. These comparisons show the general form of the eigenvalue spectra remain largely unchanged for a range of ξ_r, and the number of rattlers decreases substantially for $\xi_r > 1$ (**Figure 4D**). In addition to examining the spectra and N_r/N ratio, we also examine plots of the modes before (**Figure 5**, red) and after (**Figure 5**, blue) rattlers are removed with $\xi_r = 4$. With this choice for ξ_r, rattlers are less than 1% of the total number of agents, and we find in two separate simulation runs that large irregular vector arrows in low-m modes are no longer present. In this case, we have successfully filtered
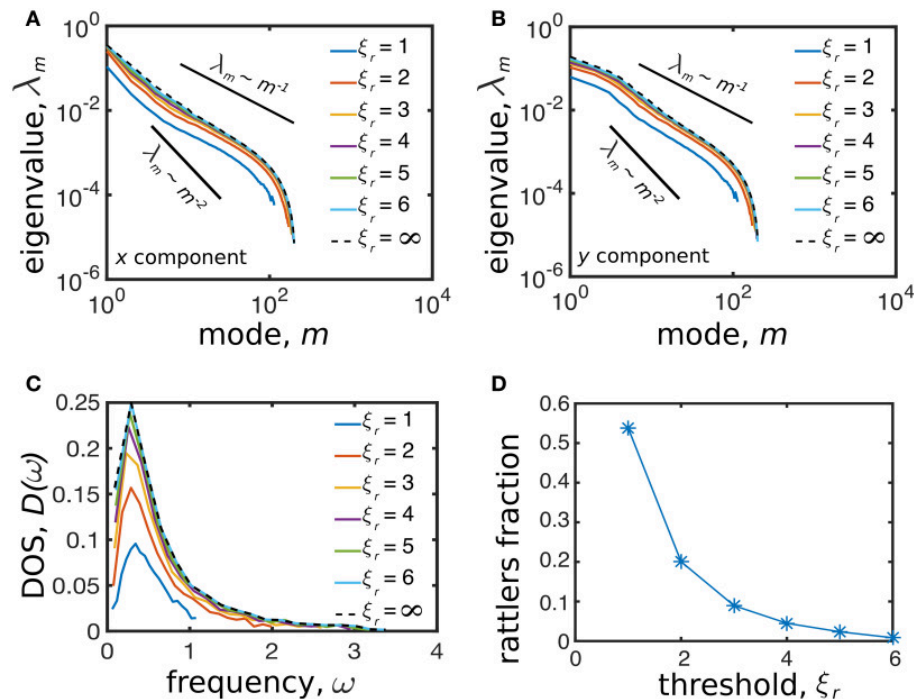
FIGURE 4 | The eigenvalue spectrum of the displacement correlation matrix before removing the rattlers (dashed black line, threshold for removal is infinitely large) and after removing the rattlers using different values of the threshold ξ_r (solid colored lines, threshold for removal is finite) for the **(A)** x and **(B)** y components. **(C)** The basic shape of the DOS is conserved for different values of the threshold ξ_r. **(D)** The fraction of rattlers monotonically decreases as the threshold value ξ_r increases.

rattler motion out from the overall collective behavior, leaving a set of $N - N_r$ eigenmodes and eigenvalues for downstream analysis.

2.2.5. (Optional) Step 5: Alternative Heuristic Method for Finding Rattlers

We briefly mention an alternative method for finding rattlers. This heuristic involves the computation of each agent's positional auto-correlation time Δt^* from the trajectory data $\vec{r}_i(t) = \langle x_i(t), y_i(t) \rangle$ and the auto-correlation functions

$$\chi_i^x(\Delta t) = \frac{1}{T} \sum_{t=1}^{T-\Delta t} \left[x_i(t) - \langle x_i \rangle \right] \times \left[x_i(t + \Delta t) - \langle x_i \rangle \right], \quad \text{and}$$

$$\chi_i^y(\Delta t) = \frac{1}{T} \sum_{t=1}^{T-\Delta t} \left[y_i(t) - \langle y_i \rangle \right] \times \left[y_i(t + \Delta t) - \langle y_i \rangle \right]. \quad (14)$$

Here, the normalizations are over all T time steps, and the averaging $\langle \cdots \rangle$ is calculated for individual agents. We seek minimal time delays Δt^x and Δt^y such that $\chi_i^x(\Delta t^x) = \chi_i^y(\Delta t^y) = 0$, which average to the auto-correlation time $\Delta t^* = (\Delta t^x + \Delta t^y)/2$. Because rattlers are free to move within their local region, they are generally not influenced by their neighbors and we expect these agents to have the highest auto-correlation times. *Critical Note 1:* The auto-correlation times are independently calculated for each of the i agents even though Δt^* does not have an explicit index i. *Critical Note 2:* Auto-correlation measurements can have non-trivial behavior that

requires individualized assessment. Some cases of note include: when multiple time delays where $\chi_i = 0$ are found, we simply take the smallest value of Δt; when there are no time delays with $\chi_i = 0$, it may be appropriate to fit a smooth function and extrapolate the delay time that χ_i intersects zero; when the auto-correlation functions are noisy and it appears that Δt as-defined is erroneous, it may be appropriate to smooth the data and infer a revised value of Δt. *Critical Note 3:* While the heuristic approach for finding rattlers has the advantage of being rapidly calculated directly from raw trajectory data, it does not have the same degree of accuracy as the threshold identification method. This cost-benefit assessment suggests the heuristic approach may find its most effective applications in real-time inference of crowd diagnostics.

Example: Testing heuristic identification of rattlers. Step 3 of the mode analysis protocol identifies rattlers based on each agent's positional fluctuations within a given mode. Here, we consider a side-by-side comparison for three simulation runs showing how a threshold of $\xi_r = 4$ compares with the distribution of auto-correlation times. Evidently, the agreement is considerable (**Figure 6**, overlapping red circles and dark squares, especially in runs 1 and 3), though certainly not perfect (**Figure 6**, non-overlapping red circles and dark squares, especially run 2). When an abundance of data and analysis time is available, the threshold approach is clearly preferable for its accuracy. However, in circumstances that limit the availability of data or when analysis is needed in near-real-time, the heuristic may be preferable.
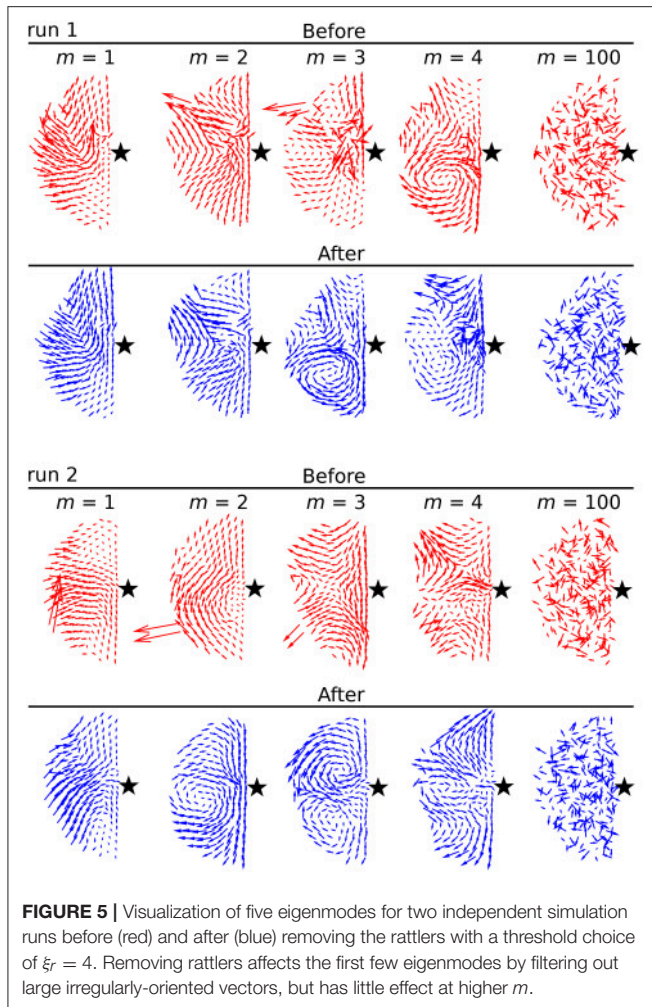
FIGURE 5 | Visualization of five eigenmodes for two independent simulation runs before (red) and after (blue) removing the rattlers with a threshold choice of $\xi_r = 4$. Removing rattlers affects the first few eigenmodes by filtering out large irregularly-oriented vectors, but has little effect at higher m.
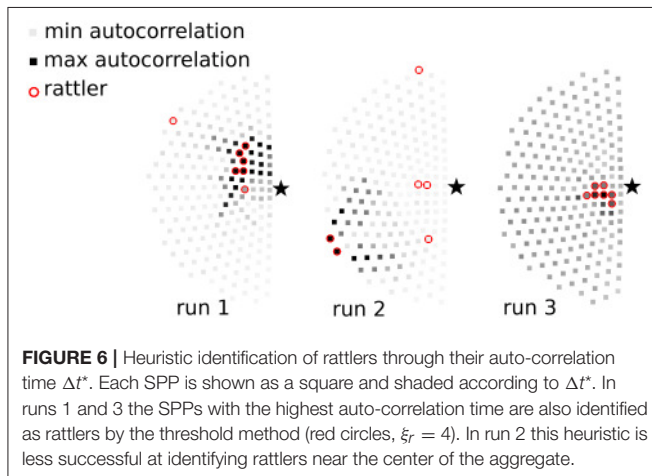
FIGURE 6 | Heuristic identification of rattlers through their auto-correlation time Δt^*. Each SPP is shown as a square and shaded according to Δt^*. In runs 1 and 3 the SPPs with the highest auto-correlation time are also identified as rattlers by the threshold method (red circles, $\xi_r = 4$). In run 2 this heuristic is less successful at identifying rattlers near the center of the aggregate.

2.2.6. (Optional) Step 6: Calculate the Density of States (DOS)

In mode analysis, the Density of States (DOS) $D(\omega^2)$ is used to quantify the probability of having a certain number of excitable oscillations at frequency ω within a given energy range $\sim \omega^2$.

This concept is well-defined at equilibrium, but more tenuous in non-equilibrium systems such as the asocial model considered here. The value of using a DOS analysis is that it sheds light on how a perturbation will transfer energy into various modes, as long as the perturbation itself does not substantially disrupt the organization of agents that defines the modal structure. To calculate $D(\omega^2)$, we note the eigenvalues are related to harmonic frequencies by $\lambda_m = \omega_m^{-2}$. Since oscillation energy is proportional to the frequency squared, $D(\omega^2)$ is essentially a histogram of the inverse-eigenvalues. *Critical Note 1:* The DOS conveys information about the rigidity of a solid; when there are many low-energy modes, the system is "soft" and will appear unstable to excitations. Here again, we mention the caveat that "low-energy modes" in the asocial model are a linearized approximation of the true response. *Critical Note 2:* A useful conceptual touchstone is the Debye law for regular lattices wherein the DOS is often expressed as the density of frequencies ω as opposed to energies ω^2. In d dimensions, the Debye law is $D(\omega) \sim \omega^{d-1}$.

Example: Interpreting eigenvalue spectra through a DOS lens. We previously analyzed the eigenvalue spectra to examine their dependence on the rattler threshold ξ_r. To provide additional context for these plots, we can explicitly plot the DOS (**Figure 4D**) to reveal the distribution of low-frequency excitations. These measurements provide a potential target of opportunity for theoretical predictions.

2.2.7. (Optional) Step 7: Find Soft Spots

When studying jammed granular materials, certain regions are often found to be partially under-constrained resulting in the presence of "soft spots" that are more likely to undergo large structural rearrangements when the system is perturbed [19]. In the context of analyzing dense human crowds, soft spots localize the people undergoing the largest displacements. The agents in these soft spots are known as "bucklers" [22], and they can be identified with a thresholding process similar to the one used to find rattlers. In this case, we seek the N_s non-rattler agents indexed by j from the collection of low-m modes with $|\vec{e}_{m,j}| \geq \langle |\vec{e}_m| \rangle + \xi_S \sigma_m$, where each term is defined as in Equations (11)–(13), and ξ_S is a yet-to-be-determined threshold for finding agents in soft spots. We also seek the N_D non-rattler agents indexed by i whose dynamics in $\langle x, y \rangle$ obey

$$\langle |\vec{r}_i(t) - \langle \vec{r}_i \rangle| \rangle \geq \sum_{i=1}^{N-N_r} \frac{\langle |\vec{r}_i(t) - \langle \vec{r}_i \rangle| \rangle}{(N - N_r)} + \xi_D \sigma_D, \qquad (15)$$

which identifies the agents whose displacement fluctuations are greater than the average by an amount equal to $\xi_D \sigma_D$. Here, σ_D is the standard deviation of displacements in $\langle x, y \rangle$ averaged over the non-rattler population, and ξ_D is another yet-to-be-determined threshold. Bucklers can now be defined as the set of agents identified by both thresholds when ξ_S and ξ_D are chosen to maximize the overlap between the two sets. This condition is quantified by the normalized agreement function (NAF)

$$\text{NAF} = \left[1 - \frac{N_S}{N - N_r} \right] \cdot \left[1 - \frac{N_D}{N - N_r} \right] \frac{|S \cap D|}{|S \cup D|}, \qquad (16)$$

where S and D are the set of agents identified by ξ_S and ξ_D, respectively. *Critical Note 1:* The bracketed prefactors in Equation (16) are weighting functions that dampen the measure of overlap if either set oversamples the total population. *Critical Note 2:* Bucklers tend to cluster in well-defined areas. In the asocial model, one of these areas is the perimeter of the SPP aggregate where the edge agents are trivially under-constrained and can be ignored as an artifact. *Critical Note 3:* An analogous thresholding can be performed with the SPP confinement pressure defined in Equation (8) and used in place of the average displacement fluctuations from Equation 15. This is useful if pressure fluctuations are hypothesized to correlate with "softness" in a specific study.

Example: Calculating thresholds and visualizing soft spots. For each of the simulation runs, we identified which agents had significant displacement fluctuations, pressure fluctuations, and mode fluctuations when the prescribed thresholds ranged from $\xi_D, \xi_P = 1$ to 7 and $\xi_S = 1$ to 5. Examining the simple overlap between these sets *without* normalization prefactors (**Figure 7**, top row), indicates a range of threshold values with substantial overlap within the sets, particularly for low threshold values. This agreement is spurious due to an oversampling of the total population at low thresholds. Using appropriate weighting functions from Equation (16) (**Figure 7**, middle row) to normalize the agreement function reveals a more well-defined range of thresholds for ξ_D and ξ_S wherein $\approx 10\%$ of the agents are in soft spots (**Figure 7**, bottom row). The agreement is maximized for $\xi_S = 2.5$ and $\xi_D = 4.5$, and we take these values to definitively identify soft spots for this example system. Other thresholds may be more appropriate for different sources of trajectory data. We also find essentially no correlation between the confinement pressure fluctuations and softness, indicating these quantities are essentially independent in the asocial model. To visualize soft spots, we plot the dense aggregate of agents and highlight individuals that were identified as having mode fluctuations above the $\xi_S = 2.5$ level in at least one of the first 10 modes (**Figure 8**). Color coding the first 10 modes for three simulation runs and recalling that agents on the perimeter are trivially under-constrained, we see soft spots are consistently localized near the core of the aggregate. This implies a high-probability of structural rearrangements will occur in this area when the system is excited.

2.2.8. (Optional) Step 8: Find Soft Modes
In the harmonic theory of crystals, eigenmodes of the displacement correlation matrix $[C_{ij}]$ fully characterize a linear response of the system to perturbations [36]. When excited, each of these "normal modes" requires an energy whose cost is inversely proportional to their corresponding eigenvalue. While useful for studying jammed granular systems, this linearized theoretical framework can break down if harmonicity or energy equipartition are violated. In such circumstances, an equivalence between the dynamical matrix C_p and displacement correlation matrix $[C_{ij}]$ becomes tenuous and deserving of

further consideration. However, information about the system's structural stability, coherent collective motion, and localized kinematics may nevertheless be conveyed through "soft modes." These modes are the eigenmodes corresponding to the highest eigenvalues of $[C_{ij}]$, which in turn, are the lowest excitation energies of the linearized theory [17–19].

To determine which modes of $[C_{ij}]$ calculated in steps 1–4 are soft, we compare the agent's eigenvalue spectrum to a spectrum arising from uncorrelated motion. Specifically, we generate a set of $N - N_r$ random displacements with zero mean and standard deviation σ_{RM}. This standard deviation of the random displacements is chosen to be equal to the standard deviation of the agent's displacements around their equilibrium positions. We then calculate the covariance matrix and eigenvalue spectrum for this Random Matrix Model (RM_σ). Comparing spectra, we now have a threshold condition: when the eigenvalues from steps 1 to 4 are *greater than* the eigenvalues from RM_σ, we identify the associated modes as *soft modes* because they relate to correlated motion; when the eigenvalues from steps 1 to 4 are *less than or equal to* the eigenvalues from RM_σ, we identify the associated modes as essentially random uncorrelated motion. *Critical Note 1:* When comparing spectra, all data should be averaged over all independent simulation runs. *Critical Note 2:* This approach has roots in principal component analysis and studies of jammed granular systems, where soft modes indicate preferential directions of relaxation in the system [16]. In terms of human crowds, we interpret soft modes as the preferential directions of collective motion because these modes feature low excitation energies, suggesting their emergence is likely to occur when a dense crowd is slightly perturbed.

Example: Using soft modes to concretely define "low-m" modes. By calculating the RM_σ spectrum and superimposing eigenvalues from simulated human crowds, we see there are up to $m = 6$ soft modes in both the x and y directions (**Figure 9A**, modes above dashed line). For practical purposes, we can now use the RM_σ threshold to quantitatively define these low-m modes as the system's soft modes. Recalling that eigenvalues and oscillation frequencies are related by $\lambda_m = \omega_m^{-2}$, we see in the DOS that correlated motions of soft modes correspond to a low-frequency Bosonic peak typically associated with long-range collective motion near the jamming transition (**Figure 9B**) [37–39]. This observation suggests the system is not mechanically stable, and that small perturbations could excite soft modes resulting in major structural rearrangements. In terms of human crowds, this would explain from a physical point of view why sudden collective motion can spontaneously emerge at high density. To visualize these collective motions, we plot the displacement vector fields for the first six soft modes from a single run (**Figure 9C**). We see they indeed carry a high degree of spatial correlation at low-m that rapidly diminishes with increasing mode number. In this example the spectra are averaged over 10 independent runs and rattlers have been removed.

3. RESULTS

The previous section described an asocial model for simulating high-density human crowds and provided a step-by-step
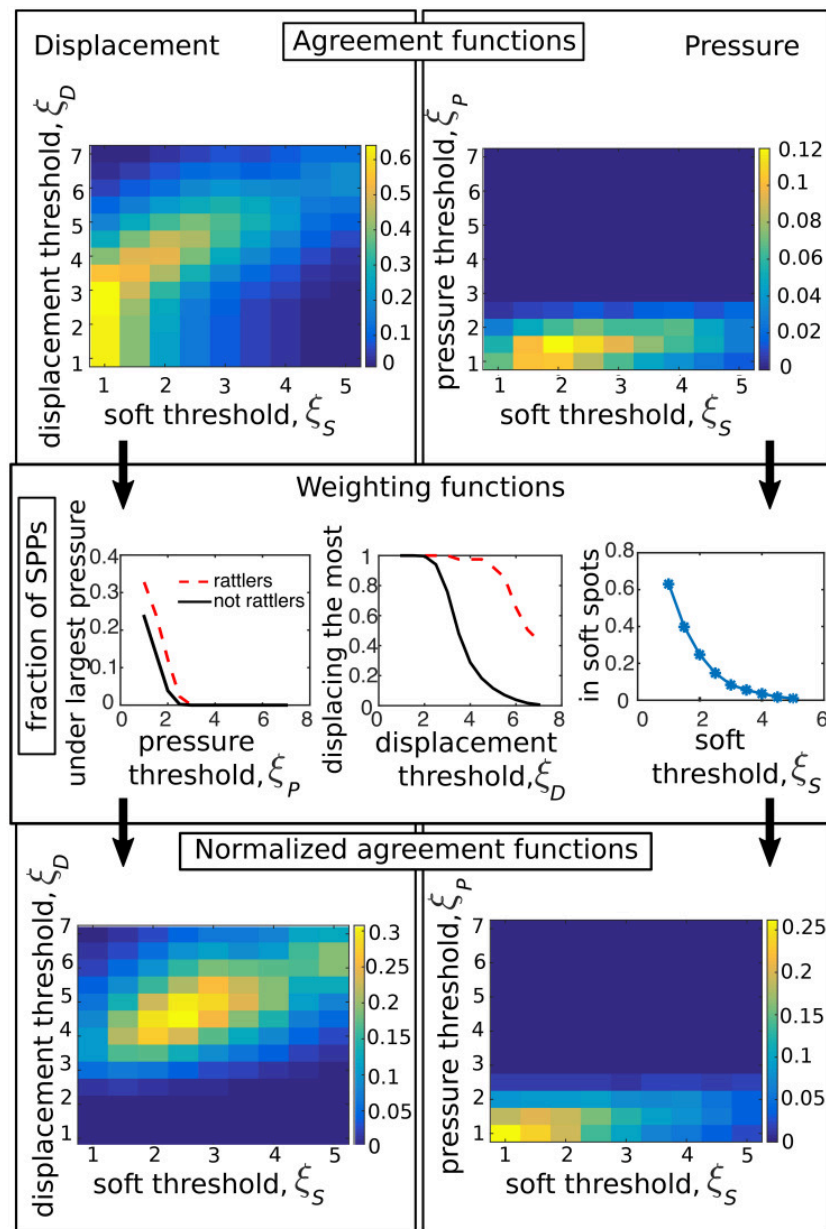
FIGURE 7 | The agreement functions for the correlation between being in a soft spot (*S*), undergoing large pressure (*P*), and large displacement (*D*) before **(top)** and after **(bottom)** the normalization by the weighting functions **(center)** that measure the size of the sets *S*, *P*, *D*. Arrows indicate the logical procedure followed to obtain the normalized agreement functions.

protocol for identifying several different types of emergent collective motion. Here, we draw inspiration from a variety of physical concepts to further interpret the information conveyed by mode analysis and its specific meaning for human collective behavior.

3.1. The $m = 1$ mode Is a Pseudo-Goldstone Mode

Symmetry plays a critical role for nearly all fields of modern physics. In condensed matter, the spontaneous breaking of

continuous symmetries is often associated with the emergence of long-range low-energy excitations known as Goldstone modes [40–42]. For example, studies of flocking in active systems have found the interaction of self-propulsion and directional alignment breaks global rotation invariance, leading to rapid collective directional changes known as the "Goldstone mode of the flock" [43–46]. In other examples, where continuous symmetries are instead explicitly broken by exogenous factors, we find pseudo-Goldstone modes, which require a small but finite energy for excitation. With these considerations in mind,

we note the asocial model involves a propulsion force aligned to a specific point of interest $\mathcal{P}$ that explicitly breaks $\langle x, y \rangle$ translational symmetry (Equation 2). We therefore predict a low-energy long-range pseudo-Goldstone mode to be a fundamental collective excitation of the asocial model.

One way to check if one or more of our system's modes is a pseudo-Goldstone mode is to examine whether the polarization correlation length is system-spanning [46]. Practically, this is accomplished by calculating the mean polarization vector $\vec{\Phi}(m) = N^{-1} \sum_{i=1}^{N} \vec{e}_m^i / |\vec{e}_m^i|$ and measuring the correlation function of each mode's fluctuations $C_m(d) = \langle [\vec{e}_m^i - \vec{\Phi}(m)] \cdot [\vec{e}_m^j - \vec{\Phi}(m)] \rangle_{d_{ij}=d}$ about this average value [46]. In this last expression, the average is over all particles i and j whose pairwise distance d_{ij} is equal to the distance d. We then define the correlation

length $l_c(m)$ as the minimum distance at which $C_m(l_c) = 0$. Plotting $C_m(d)$ for a few modes shows that most have a relatively short correlation length, while the $m = 1$ mode extends across the entire system (**Figure 10**). This system-wide excitation is at the lowest possible mode number, which in linear mode theory corresponds to the lowest possible excitation energy. Thus, these two pieces of evidence combine to strongly implicate the $m = 1$ mode as a pseudo-Goldstone mode. Because its origins can be traced to a broken continuous symmetry, this long-range highly correlated collective motion is an intrinsic effect of densely aggregated agents.

In our asocial model, the $m = 1$ mode is a collective motion that slides up-and-down along the right-most edge of the simulation box (**Figure 9C**). In real-world circumstances this means the most easily-excitable mode would result in a large number of people being suddenly displaced together, possibly toward a wall, concert stage, or some other barrier. Such a situation has been widely observed in conjunction with the emergence of shock waves and density waves during "crowd turbulence" [12, 47]. Since its origins can be traced to the general principle of symmetry breaking, this type of long-range collective motion can be expected as a latent excitation arising in a wide range of circumstances with the potential for causing crowd crush casualties [48].

3.2. Topological Defect Density Drives Disorder in the Modes

If broken symmetries and pseudo-Goldstone modes can explain the $m = 1$ long-range collective motion, then what is the most useful way to understand the remaining $m > 1$ modes?
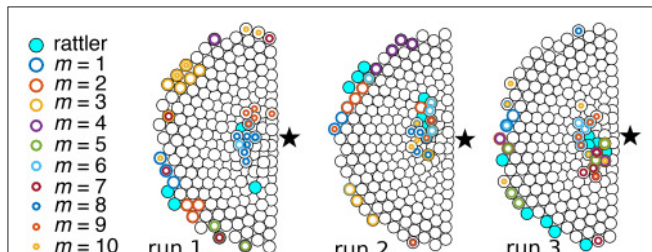
FIGURE 8 | Identification of soft spots for three independent simulation runs. Each SPP is shown as a disk and the star ($\star$) represents the point of interest $\mathcal{P}$. Superimposing data for the first 10 modes shows that soft spots appear in multiple modes in the same general area near the core of the aggregate. Apparent soft spots along the periphery are artifacts due to under-constrained edge effects.

FIGURE 9 | Eigenmode analysis of asocial model for high-density human crowds. **(A)** Eigenvalue spectrum λ_m of the displacement correlation matrix exhibits scaling properties between $\lambda_m \sim m^{-1}$ and $\sim m^{-2}$ (black solid lines). Eigenmodes up to $m = 6$ in both x (blue) and y (orange) directions are larger than the random matrix model (RM_σ, dashed line), thus they are named "soft modes" and describe correlated motion. **(B)** The DOS exhibits a Bosonic peak in both the x and y components, indicating mechanical instability. **(C)** Soft mode vector fields for run 1 ($m = 1$ to 6) are more spatially correlated than a mode below RM_σ ($m = 15$).
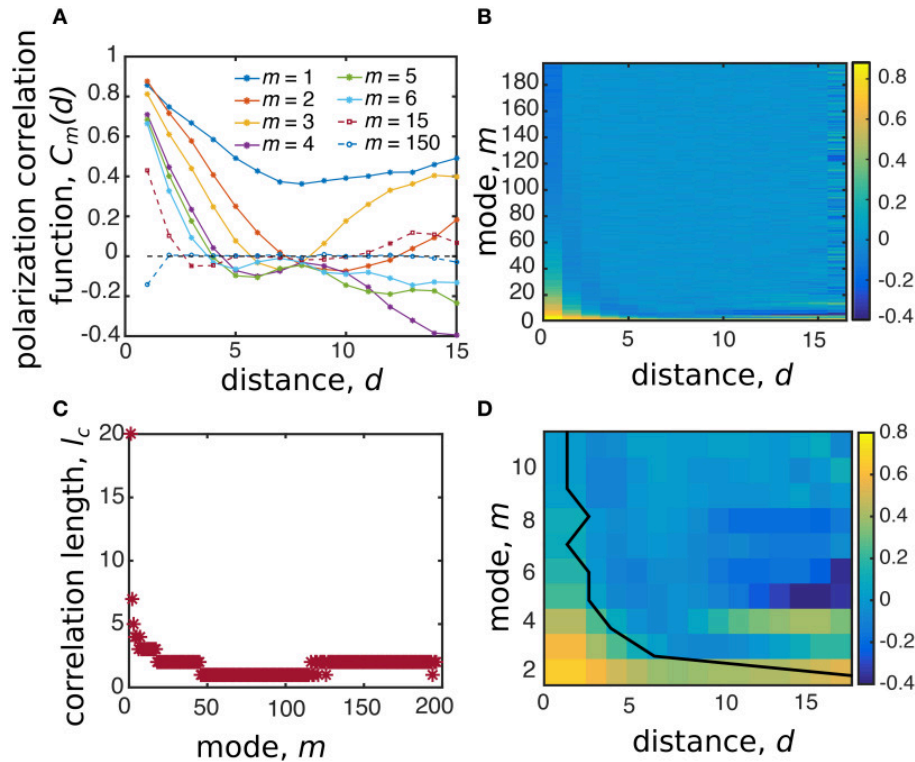
FIGURE 10 | The polarization correlation length identifies the $m = 1$ mode as a pseudo-Goldstone mode. **(A)** The correlation function $C_m(d)$ for the mode polarization fluctuations plotted against distance between particles d decays rapidly for modes $m > 1$, which can be seen in a heat map **(B)** for all modes. **(C)** The correlation length for all modes defined as the distance ℓ_C where $C_m(\ell_C) = 0$. **(D)** Examining the first 12 modes as a heat map and superimposing the correlation length number illustrates where $C_m(d) = 0$.

Remarkably, and somewhat surprisingly, topological principles provide useful insights. Two modes are considered topologically equivalent if their vector fields can be continuously deformed to match one another, and as such, the difference in excitation energy will become arbitrarily small as their alignment converges [42]. However, the introduction of a topological defect, such as vortices in superfluids, magnetic flux tubes in superconductors and edge-dislocations in liquid crystals, prevent convergence and drive a persistent non-zero energetic difference. To check whether topological defects play a meaningful role in explaining the structure of eigenmode m, we calculate the winding number charge $q_i = (2\pi)^{-1} \oint_\ell \vec{\nabla}\theta \cdot \vec{d\ell_i}$ for each non-edge agent i using a path ℓ_i that loops around i's nearest-neighbors. Here, $\vec{\nabla}\theta$ measures the change in orientation between each agent's eigenmode vector along the loop. This measure is 0 when there are no topological defects, $+1$ if there is a vortex centered on the agent, or -1 if there is an anti-vortex centered on the agent. We can therefore identify each agent as coinciding with a topological defect depending on whether the local vector field makes a vortex with positive or negative charge q_i. Qualitatively examining q_i for a handful of different modes, it is clear that there are a number of topological defects, especially for $m > 1$ (**Figure 11A**). Recognizing that the rotation of the mode's vector field around these defects will reduce the correlation length, we sum the total absolute defect charge $Q_{\text{abs}}(m) = \sum_{i=1}^{N} |q_i|$ for each

mode (**Figure 11B**) and estimate the expected correlation length. Assuming n defects are randomly distributed among N agents in a half disk of radius R and area $\pi R^2/2$, the spatial distribution of defects can be expressed by a Poisson process of intensity ρ_q:

$$P(n, R) = \frac{[\rho_q \pi R^2]^n}{2n!} e^{-\rho_q \pi R^2/2}, \quad (17)$$

where $\rho_q = Q_{\text{abs}}/N$ is the defect density. The probability to find the first defect closest to the disk's origin at a distance greater than R is equal to $F(> R) = P(\text{no point within } R) = \exp[-\rho_q \pi R^2/2]$. Differentiating with respect to R and being careful with the sign of $F(> R)$, we find the probability $f(R)$ for the first neighbor at a distance R is $f(R) = \pi \rho_q R e^{-\rho_q \pi R^2/2}$. Thus, the average nearest-neighbor distance $\langle R \rangle$ between any two points in the semi-disc is:

$$\langle R \rangle = \int_0^\infty R f(R) dR,$$
$$= \int_0^\infty \pi \rho_q R^2 e^{-\rho_q \pi R^2/2} dR,$$
$$= -\left[R e^{-\rho_q \pi R^2/2} \right]_0^\infty + \int_0^\infty e^{-\rho_q \pi R^2/2} dR,$$

FIGURE 11 | Topological defects drive disorder in $m > 1$ modes. **(A)** Winding number charge q_i for the soft modes of two independent simulation runs. Each SPP is colored according to its winding number and topological defects are centered on SPPs with $q_i \neq 0$. In both simulation runs the number of topological defects increases with mode number. **(B)** The total absolute charge Q_{abs} plotted against mode number m for 10 independent simulation runs shows the number of defects increases with increasing mode number. **(C)** Scatterplot of the total absolute charge Q_{abs} vs. correlation length l_c for each mode of the 10 independent runs (circles) compared to the parameter-free predicted estimate from our geometric argument (line).

$$
= \sqrt{\frac{2}{\pi \rho_q}} \int_0^\infty e^{-t^2} \, dt,
$$

$$
= \frac{1}{\sqrt{2\rho_q}}, \tag{18}
$$

which is the average distance between two topological defects at density ρ_q. To connect $\langle R \rangle$ with the polarization fluctuation correlation length, we notice that if two neighboring defects are both positive or both negative, the mode's vector field cannot change direction in the space separating them, therefore $l_c = \langle R \rangle$. However, if two neighboring defects have opposite signs, the vector field must change sign and $l_c \approx \langle R \rangle / 2$. Therefore, on average we have

$$
\langle l_c(m) \rangle = \frac{1}{2} \langle R \rangle + \frac{1}{2} \frac{\langle R \rangle}{2} = \frac{3}{4} \sqrt{\frac{N}{2 Q_{abs}(m)}}. \tag{19}
$$

This last expression is a parameter-free theoretical prediction that readily agrees with our numerical computations (**Figure 11C**),

suggesting that disordered features of modes $m > 1$ arise from topological defects distributed throughout the system. For $m < 50$ (**Figure 11B**), we interpret this result as indicating that modes cannot be continuously deformed into lower energy collective excitations, while for $m > 50$ a maximum disorder is reached due to a saturation in defect density.

3.3. Collective Motion has Microscopic Structural Origins

Understanding collective motion in high-density crowds is motivated by an impetus to predict and prevent human disaster. Thus far, we have successfully linked individual trajectories to emergent collective motion through mode analysis. While the analysis provides insights such as the *existence* of soft spots and psuedo-Goldstone excitations, the specific *locations and orientations* of these phenomena depend on trajectory data. Consequently, these details can only be unlocked through an *after the fact* analysis, as opposed to the more-desirable goal of assessing risk in *real-time*. Nevertheless, it should still be possible to perform real-time risk assessment given that these collective motions ultimately depend on the microscopic self-organized structure of the crowd. The question remains: how do we make such an inference?

When we examine the structure of high-density crowds, it clearly deviates from the well-known hexagonal packing of 2D hard disks under uniform conditions. In the asocial model, these irregularities arise from a pressure gradient, which is calculated by averaging Equation (8) as a function of distance from the point of interest $\mathcal{P}$ (**Figure 12A**). The structural irregularities created by this pressure gradient can be quantified by similarly computing the *average* number of interacting neighbors, also as a function of distance from $\mathcal{P}$. Within the bulk of the crowd, this value is typically equal to six (**Figure 12B**, black line), which would be expected for homogeneous packing, while increasing to seven near $\mathcal{P}$ due to the higher pressure (**Figure 12B**, dashed black line). If we now filter our averaging to examine just the bucklers within soft spots, we see the average number of interacting neighbors is measurably *higher*, suggesting that the local coordination number contains a signature of these potentially high-risk locations (**Figure 12B**, solid red line). Examining specific runs and comparing the *local* coordination number (**Figure 12C**) to the location of bucklers in soft spots (**Figure 12D**), we find a broad consistency between these two measures. The critical point here is that the coordination number can be extracted from a single "snap-shot" of the crowd, whereas soft spots are identified through the full machinery of mode analysis. Even if the mapping is not perfectly one-to-one, coordination number may provide a valuable correlate for predicting the location of high-risk areas before collective motions become excited.

If soft spots are indeed a consequence of local structure, the mechanistic connection remains to be identified and understood. Therefore, we measure the two particle radial structure factor $g(d) = [(N - N_r)(N - N_r - 1)]^{-1} \sum_{i=1}^{N-N_r} \sum_{j \neq i}^{N-N_r} \delta(d - d_{ij})$, which quantifies the radial distribution of distances between neighboring agents. Unlike globally averaged properties, $g(d)$

FIGURE 12 | Local coordination number correlates with soft spots. **(A)** The average radial pressure normalized by $P_0 = \mu v_0 (2\pi r_0)^{-1}$ reveals a gradient within the aggregate. This pressure drives deviations **(B)** from homogenous hexagonal packing as measured by the average number of interacting neighbors at a distance r from the point of interest $\mathcal{P}$. Hexagonal close packings of hard disks have six interacting neighbors, which is added as a dashed line for reference. **(C)** The local coordination number, shown here for 4 runs, can be computed from a snap shot of the system. Comparison with **(D)** shows the coordination number strongly correlates with the location of soft spots, which are found through more lengthy computations involving the identification of aggregated bucklers (red).

FIGURE 13 | The structure factor $g(d)$ helps explain why local coordination predicts location of soft spots. **(A)** The structure factor $g(d)$ shows short range order (peaks for $d \lesssim 3$), but generally no long-range order. **(B)** Zooming in on the blue boxed region from **(A)**, we see differences in local structure between bucklers in soft spots (red solid line) and the averaged aggregate (black dashed line). All data is generating by averaging over 10 independent simulation runs.

provides information about the local structure [38]. For the asocial model, it reveals that the overall structure has clear short range order (**Figure 13**, peaks for $d < 3$) but no periodic long range order (**Figure 13**, generally smooth distribution for $d > 3$). We see the position of the first peak centered at $d = 0.8$ suggests that, on average, the agents slightly overlap due to the self-propulsion forces. When we filter our measurement and reexamine $g(d)$ strictly for the bucklers in soft spots, the data shows a new sub-peak around $0.5 \lesssim d \lesssim 0.8$, while a second more prominent peak shifts to $d \approx 0.9$. This seems to suggest the structure within a soft spot is asymmetrically squeezed with nearest-neighbors somewhat closer than average in one direction, presumably in the direction of $\mathcal{P}$, while also somewhat further away than average from other neighbors. As a result, this irregular structure provides a microscopic mechanism for bucklers to easily displace when perturbed.

In terms of high-density human crowds, motion can be thought of as the superposition of the most easily excited
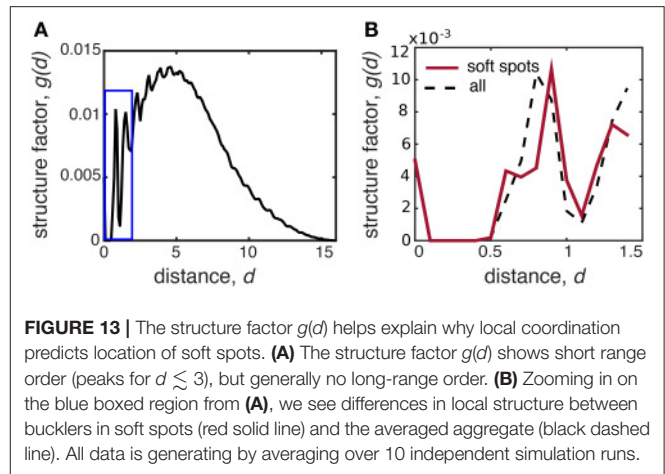
modes. When motion occurs, our analysis predicts that people in soft spots would be the ones displacing the most. We therefore interpret soft spots as posing the highest risk for tripping and subsequent trampling, especially if activated by a sudden and unexpected external perturbation. Qualitatively, this phenomenon has been observed in a number of crowd disasters, when sudden unexpected movements of the crowd cause individuals to trip and fall, resulting in injury or death due to trampling or compressive asphyxia [12, 48, 49]. Furthermore, the observation that these areas can be heuristically detected through heterogeneity in the local coordination number provides a potential target for real-time prediction and prevention.

3.4. The "Participation Ratio" and "Effective Coordination Number" Do Not Measure Collectiveness in the Asocial Model

While we were able to successfully co-localize soft spots with the local coordination number, there are two additional metrics we tested that were found to provide insufficiently detailed information. Nevertheless, because these metrics are more widely used to study densely-packed jammed systems, we provide an overview of our findings so that others may have our null results as a reference. Our main finding with these metrics is that they seem to detect a difference between high- and low-m modes with a transition around $m \approx 50$, similar to the transition found when measuring the density of topological defects (**Figure 11B**). However, this says little about soft modes, the eigenvalue spectrum, the DOS, or auto-correlation length. Any deeper significance for active matter systems apparently requires further analysis.

A standard measure for the collectiveness of a mode is given by the participation ratio $\mathrm{PR}(m)$, which quantifies how many agents in the system move when a given mode is excited. In the literature there are several definitions of this metric [16, 21], and while we tested them all, we present results when $\mathrm{PR}(m)$ is calculated as

$$\mathrm{PR}(m) = \frac{\left(\sum_{i=1}^{N} |\vec{e}_m^i|^2 \right)^2}{N \sum_{i=1}^{N} |\vec{e}_m^i|^4}, \tag{20}$$

FIGURE 14 | While consistent with our general findings, not all measures of structure explain collective motion. **(A)** The participation ratio PR(m) initially decreases as a function of mode number, but after $m \sim 6$ increases toward the random case (dotted line). **(B)** The effective coordination number $z_{\text{eff}}(m)$ shows the lowest modes are under-constrained.

4. DISCUSSION

With an eye toward understanding, predicting, and preventing tragedies at mass gatherings, we view our main results as revealing mechanisms for the emergence of potentially dangerous collective motion. By first identifying these principles and outlining a quantitative framework for measuring their existence, we are now in position to test their real-world applicability using video data of concerts, pilgrimages, and sporting events. This next step is a straight-forward empirical data collection process, given the current availability of low-cost high-definition digital cameras and inexpensive cloud-computing resources for rapid image analysis. The only remaining obstacle, therefore, is to develop computer vision algorithms that robustly and automatically track individual trajectories in footage of high-density crowds. While this image analysis challenge is open-ended, it may be sufficient for our purposes to simply study coarse-grained fields of view that average motion over regional domains encompassing several people.

If the methods outlined here prove to be broadly predictive in describing high-density human collective motion when no disasters occur, then they will become a valuable starting point for developing conceptually new strategies that enhance safety at mass gatherings. In the long term, we hope our results will lead to practical tools for real-time monitoring and predictive diagnostics at mass events. We also note that while the techniques described here are motivated by human crowds, they provide an analytical framework for extracting key insights from other real-world problems such as the characterization of biological tissues, the dynamics of migrating cancer cells, animal collective motion, real-time material characterization, and self-monitoring industrial assembly-lines.

which, respectively, takes values between 0 and 1 for fully localized and fully extended collective motion [21]. Plotting the participation ratio against mode number m provides a signature of the system and gives an overview of the collective nature of modes. In our case, we find the participation ratio for soft modes is lower than the random matrix RM_σ, and increases toward 1/2 with mode number m (**Figure 14A**). This occurs because the typical length of the displacements on the high-m modes are highly similar while their direction is random. Conversely, soft mode displacements are more variable in length but more correlated in direction. This seems to suggest that in the framework considered here, the participation ratio is not an appropriate measure for detecting collective behavior.

Another commonly used metric to characterize modes is the effective coordination number [21]:

$$z_{\text{eff}}(m) = \frac{\left(\sum_{i=1}^{N} z_i |\vec{e}_i^m|^2 \right)}{\sum_{i=1}^{N} |\vec{e}_i^m|^2} - 3, \quad (21)$$

where z_i is the number of neighbors interacting with the i^{th} agent defined as $d_{ij} < 2r_0$, and -3 is to remove degrees of freedom associated with global 2D rotational/translational symmetries. This expression calculates the average number of constraints per agent in each mode, weighted by their displacement on that mode. In jammed solids, its value depends on the amount of compression and affects the frequency ω of modes [21]. Here, we cannot precisely relate the eigenvalues to the energy of the system, therefore $z_{\text{eff}}(m)$ simply helps identify over- or under-constrained modes. In particular, rigid stability requires $z_{\text{eff}}(m) \geq 3$; in light of our results (**Figure 14B**), the system appears generally non-rigid.

AUTHOR CONTRIBUTIONS

FUNDING

ACKNOWLEDGMENTS

REFERENCES

1. Holland JH. *Complexity: A Very Short Introduction.* Oxford: Oxford University Press (2014).
2. Thaler RH, Sunstein CR. *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Westminster: Penguin Books (2009).
3. Thaler RH. *Misbehaving: The Making of Behavioral Economics.* New York, NY: W. W. Norton and Company (2016).
4. Hausman DM, Welch B. Debate: to nudge or not to nudge. *J Polit Philos.* (2010) **18**:123–36. doi: 10.1111/j.1467-9760.2009.00351.x
5. King R, Churchill EF, Tan C. *Designing with Data: Improving User Experience with A/B Testing.* Sebastopol, CA: O'Reilly Media (2017).

6. Kohavi R, Thomke S. The surprising power of online experiments. *Harvard Bus Rev.* (2017) **5**:74–82. https://hbr.org/2017/09/the-surprising-power-of-online-experiments

7. Aven T. On the meaning of a black swan in a risk context. *Safety Sci.* (2013) **57**:44–51. doi: 10.1016/j.ssci.2013.01.016

8. Taleb NN. *The Black Swan: The Impact of the Highly Improbable.* Random House Trade Paperbacks (2010).

9. Milsten AM, Maguire BJ, Bissell RA, Seaman KG. Mass-gathering medical care: a review of the literature. *Prehosp Disast Med.* (2002) **17**:151–62. doi: 10.1017/S1049023X00000388

10. Carley S, Mackway-Jones K, Donnan S. Major incidents in Britain over the past 28 years: The case for the centralised reporting of major incidents. *J Epidemiol Commun Health* (1998) **52**:392–8. doi: 10.1136/jech.52.6.392

11. Wagner U, Fälker A, Wenzel V. Fatal incidents by crowd crush during mass events:(Un) preventable phenomenon. *Anaesthesist Der.* (2013) **62**:39–46. doi: 10.1007/s00101-012-2124-z

12. Helbing D, Johansson A, Al-Abideen HZ. Dynamics of crowd disasters: an empirical study. *Phys Rev E* (2007) **75**:046109. doi: 10.1103/PhysRevE.75.046109

13. Helbing D, Johansson A, Mathiesen J, Jensen MH, Hansen A. Analytical approach to continuous and intermittent bottleneck flows. *Phys Rev Lett.* (2006) **97**:168001. doi: 10.1103/PhysRevLett.97.168001

14. Cristiani E, Piccoli B, Tosin A. Multiscale modeling of granular flows with application to crowd dynamics. *Multiscale Model Sim.* (2011) **9**:155–82. doi: 10.1137/100797515

15. Faure S, Maury B. Crowd motion from the granular standpoint. *Math Mod Meth Appl S* (2015) **25**:463–93. doi: 10.1142/S0218202515400035

16. Henkes S, Brito C, Dauchot O. Extracting vibrational modes from fluctuations: a pedagogical discussion. *Soft Matter* (2012) **8**:6092–109. doi: 10.1039/c2sm07445a

17. Brito C, Dauchot O, Biroli G, Bouchaud JP. Elementary excitation modes in a granular glass above jamming. *Soft Matter* (2010) **6**:3013–22. doi: 10.1039/c001360a

18. Ashton DJ, Garrahan JP. Relationship between vibrations and dynamical heterogeneity in a model glass former: extended soft modes but local relaxation. *Eur Phys J E* (2009) **30**:303–7. doi: 10.1140/epje/i2009-10531-6

19. Manning ML, Liu AJ. Vibrational modes identify soft spots in a sheared disordered packing. *Phys Rev Lett.* (2011) **107**:108302. doi: 10.1103/PhysRevLett.107.108302

20. Henkes S, Fily Y, Marchetti MC. Active jamming: self-propelled soft particles at high density. *Phys Rev E* (2011) **84**:040301. doi: 10.1103/PhysRevE.84.040301

21. Xu N, Vitelli V, Liu A, Nagel S. Anharmonic and quasi-localized vibrations in jammed solids: modes for mechanical failure. *Europhys Lett.* (2010) **90**:56001. doi: 10.1209/0295-5075/90/56001

22. Charbonneau P, Corwin EI, Parisi G, Zamponi F. Jamming criticality revealed by removing localized buckling excitations. *Phys Rev Lett.* (2015) **114**:125504. doi: 10.1103/PhysRevLett.114.125504

23. Bottinelli A, Sumpter DTJ, Silverberg JL. Emergent structural mechanisms for high-density collective motion inspired by human crowds. *Phys Rev Lett.* (2016) **117**:228301. doi: 10.1103/PhysRevLett.117.228301

24. Silverberg JL, Bierbaum M, Sethna JP, Cohen I. Collective motion of humans in mosh and circle pits at heavy metal concerts. *Phys Rev Lett.* (2013) **110**:228701. doi: 10.1103/PhysRevLett.110.228701

25. Vicsek T, Zafeiris A. Collective motion. *Phys Rep.* (2012) **517**:71–140. doi: 10.1016/j.physrep.2012.03.004

26. Sumpter DJ. *Collective Animal Behavior.* Princeton, NJ: Princeton University Press (2010).

27. Helbing D, Molnar P. Social force model for pedestrian dynamics. *Phys Rev E.* (1995) **51**:4282. doi: 10.1103/PhysRevE.51.4282

28. Helbing D. A mathematical model for the behavior of pedestrians. *Behav Sci.* (1991) **36**:298–310. doi: 10.1002/bs.3830360405

29. Vivek S, Yanni D, Yunker PJ, Silverberg JL. Anticipating and inoculating against hacks that exploit the emergent collective motion of autonomous vehicles. *arXiv:170803791* (2017).

30. Helbing D, Farkas I, Vicsek T. Simulating dynamical features of escape panic. *Nature* (2000) **407**:487–90. doi: 10.1038/35035023

31. Moussaïd M, Helbing D, Theraulaz G. How simple rules determine pedestrian behavior and crowd disasters. *Proc Natl Acad Sci USA.* (2011) **108**:6884–8. doi: 10.1073/pnas.1016507108

32. Helbing D, Johansson A. Pedestrian, crowd, and evacuation dynamics. *arXiv:13091609* (2013).

33. Landau LD, Lifshitz E. *Course of Theoretical Physics: Theory of Elasticity,* Vol. 7. New York, NY: Elsevier (1986).

34. Van Hecke M. Jamming of soft particles: geometry, mechanics, scaling and isostaticity. *J Phys Condens Matter* (2009) **22**:033101. doi: 10.1088/0953-8984/22/3/033101

35. Majmudar T, Sperl M, Luding S, Behringer RP. Jamming transition in granular systems. *Phys Rev Lett.* (2007) **98**:058001. doi: 10.1103/PhysRevLett.98.058001

36. Ashcroft NW, Mermin ND. *Solid State Physics.* Stamford, CT: Thomson Learning (1976).

37. Wyart M, Nagel S, Witten T. Geometric origin of excess low-frequency vibrational modes in weakly connected amorphous solids. *Europhys Lett.* (2005) **72**:486. doi: 10.1209/epl/i2005-10245-5

38. Laird BB, Schober H. Localized low-frequency vibrational modes in a simple model glass. *Phys Rev Lett.* (1991) **66**:636. doi: 10.1103/PhysRevLett. 66.636

39. Silbert LE, Liu AJ, Nagel SR. Vibrations and diverging length scales near the unjamming transition. *Phys Rev Lett.* (2005) **95**:098301. doi: 10.1103/PhysRevLett.95.098301

40. Nambu Y. Quasi-particles and gauge invariance in the theory of superconductivity. *Phys Rev.* (1960) **117**:648. doi: 10.1103/PhysRev.117.648

41. Goldstone J. Field theories with superconductor solutions. *Il Nuovo Cimento* (1961) **19**:154–64.

42. Sethna J. *Statistical Mechanics: Entropy, Order Parameters, and Complexity.* Oxford, UK: Oxford University Press (2006).

43. Toner J, Tu Y, Ramaswamy S. Hydrodynamics and phases of flocks. *Ann Phys.* (2005) **318**:170–244. doi: 10.1016/j.aop.2005.04.011

44. Ramaswamy S. The mechanics and statistics of active matter. *Annu Rev Condens Matter Phys.* (2010) **1**:323–45. doi: 10.1146/annurev-conmatphys-070909-104101

45. Cavagna A, Cimarelli A, Giardina I, Parisi G, Santagati R, Stefanini F, et al. Scale-free correlations in starling flocks. *Proc Natl Acad Sci USA* (2010) **107**:11865–70. doi: 10.1073/pnas.1005766107

46. Bialek W, Cavagna A, Giardina I, Mora T, Pohl O, Silvestri E, et al. Social interactions dominate speed control in poising natural flocks near criticality. *Proc Natl Acad Sci USA* (2014) **111**:7212–7. doi: 10.1073/pnas.13240 45111

47. Takatori S, Yan W, Brady J. Swim pressure: stress generation in active matter. *Phys Rev Lett.* (2014) **113**:028103. doi: 10.1103/PhysRevLett.113. 028103

48. Fruin JJ. The causes and prevention of crowd disasters. *Eng Crowd Safety* (1993) **1**:10.

49. Krausz B, Bauckhage C. Loveparade 2010: automatic video analysis of a crowd disaster. *Comput Vis Image Und.* (2012) **116**:307–19. doi: 10.1016/j.cviu.2011.08.006

Insights into the Social Behavior of Surface and Cave-Dwelling Fish (*Poecilia mexicana*) in Light and Darkness through the Use of a Biomimetic Robot

David Bierbach[1,2]*, Juliane Lukas[1,2,3], Anja Bergmann[1], Kristiane Elsner[1], Leander Höhne[1], Christiane Weber[1], Nils Weimar[1], Lenin Arias-Rodriguez[4], Hauke J. Mönck[5], Hai Nguyen[2], Pawel Romanczuk[6], Tim Landgraf[5] and Jens Krause[2,3]

[1] Humboldt-Universität zu Berlin, bologna.lab, Q-Team Programm, Berlin, Germany, [2] Department of Biology and Ecology of Fishes, Leibniz-Institute of Freshwater Ecology and Inland Fisheries, Berlin, Germany, [3] Faculty of Life Sciences, Thaer Institute, Humboldt University of Berlin, Berlin, Germany, [4] División Académica de Ciencias Biológicas, Universidad Juárez Autónoma de Tabasco (UJAT), Villahermosa, Tabasco, Mexico, [5] Freie Universität Berlin, FB Mathematik u. Informatik, Berlin, Germany, [6] Department of Biology, Institute for Theoretical Biology, Humboldt Universität zu Berlin, Berlin, Germany

Biomimetic robots (BRs) are becoming more common in behavioral research and, if they are accepted as conspecifics, allow for new forms of experimental manipulations of social interactions. Nevertheless, it is often not clear which cues emanating from a BR are actually used as communicative signals and how species or populations with different sensory makeups react to specific types of BRs. We herein present results from experiments using two populations of livebearing fishes that differ in their sensory capabilities. In the South of Mexico, surface-dwelling mollies (*Poecilia mexicana*) successfully invaded caves and adapted to dark conditions. While almost without pigment, these cave mollies possess smaller but still functional eyes. Although previous studies found cave mollies to show reduced shoaling preferences with conspecifics in light compared to surface mollies, it is assumed that they possess specialized adaptations to maintain some kind of sociality also in their dark habitats. By testing surface- and cave-dwelling mollies with RoboFish, a BR made for use in laboratory experiments with guppies and sticklebacks, we asked to what extent visual and non-visual cues play a role in their social behavior. Both cave- and surface-dwelling mollies followed the BR as well as a live companion when tested in light. However, when tested in darkness, only surface-dwelling fish were attracted by a live conspecific, whereas cave-dwelling fish were not. Neither cave- nor surface-dwelling mollies were attracted to RoboFish in darkness. This is the first study to use BRs for the investigation of social behavior in mollies and to compare responses to BRs both in light and darkness. As our RoboFish is accepted as conspecific by both used populations of the Atlantic molly only under light conditions but not in darkness, we argue that our replica is providing mostly visual cues.

Keywords: RoboFish, *Poecilia mexicana*, cave molly, Atlantic molly, biomimetic robot

INTRODUCTION

Biomimetic robots (BRs) are becoming more common in behavioral research (Webb, 2000; Krause et al., 2011; Butail et al., 2015). One of the major advantages of BRs is that social interactions that are often characterized by mutual influences (Herbert-Read et al., 2012; Jolles et al., 2017) and feedbacks between multiple individuals (Harcourt et al., 2009) become in part controllable by the experimenter (Krause et al., 2011). Thus, standardized testing and new forms of experimental manipulations of social interactions can be achieved using BRs as interaction partners. However, investigations of social behavior become only meaningful when live animals accept BRs as conspecifics [(Landgraf et al., 2016), see similar views for computer animations in behavioral ecology (Chouinard-Thuly et al., 2017)]. At the moment, the number of animal species that respond to BRs as conspecifics is quite small. Thus, it is urgently needed to explore which cues emanating from BRs are crucial for the acceptance as a conspecific [reviewed in Landgraf et al. (2016)]. Tinbergen (1948) proposed that only a small subset of perceivable cues are actual signals ("social releasers"). They can be species specific, and animals often use sets of multiple cues to assess their (social) environment (Candolin, 2003). The identification of relevant cues and their realistic imitation is one of the most challenging parts in developing BRs (Krause et al., 2011). Since different species possess different sensory capabilities (Burnett, 2011), comparing the response of species with known differences in sensory ecologies toward BRs might help developers to understand the cues that are most important to establish social acceptance of their respective BRs. As an obvious by-product, BRs can help researchers to gain a much better understanding of social interactions in species, population, or ecotypes with different ecological and evolutionary backgrounds. In summary, both developers and biologists using BRs will benefit from a broader list of species investigated with BRs.

Here, we explored whether two populations of the Atlantic molly (*Poecilia mexicana*, a cave-dwelling and a surface-dwelling ecotype) accept a BR as a conspecific that was initially developed to meet the requirements of the closely related guppy (*Poecilia reticulata*), a "model organism" in many different biological fields (Magurran, 2005). Both populations of cave- and surface-dwelling mollies differ in their evolutionary background and consequently also in their assumed sensory capabilities and social tendencies. Although cave mollies still possess functional albeit smaller eyes than their surface-dwelling counterparts (Körner et al., 2006; Eifert et al., 2014), their non-visual systems seem to be much better developed (Parzefall, 1969, 1970, 2001; Peters et al., 1973; Parzefall et al., 2007) with recent investigations pointing out that both chemical and mechanosensory communication is more pronounced in cave mollies compared to surface dwellers (Rüschenbaum and Schlupp, 2013; Jourdan et al., 2016). Furthermore, cave mollies were previously found to be less attracted to conspecifics in dichotomous choice tests under normal light conditions and hence are assumed to be less social than the closely related surface ecotypes (Plath and Schlupp, 2008).

We tested the social behavior of surface- and cave-dwelling Atlantic mollies with both a live conspecific and a BR under both light and dark conditions. We hypothesized that cave mollies should be generally less social toward (e.g., do not follow closely) live conspecifics and BRs compared to surface-dwelling mollies in light but should be better able to maintain some sociality in darkness with both a live conspecific and a BR compared to surface-dwelling fish. Our study not only tests for differences in social behavior of surface and cave-dwelling fish but also tests whether BRs constructed as mobile replicas are accepted as conspecifics when visual cues are omitted (as in darkness).

MATERIALS AND METHODS

Test Fish and Their Maintenance

In this study, we used second-generation lab-reared descendants of wild-caught fish from two populations of the Atlantic molly (*P. mexicana*) that were caught during field trips to the Tacotalpa river system in Tabasco, Mexico (Tobler et al., 2011; Plath et al., 2013). Our surface population originated from the Río Oxolotán, a tributary to the Río Grijalva, while our cave population stemmed from chamber 7 of the cave Cueva del Azufres (**Figure 1**). Fish were reared in randomly outbred mixed-sex stocks at the Laboratory of Genetics and Ecophysiology from the Academic Division for Biological Sciences-UJAT. Some of them were transported to the Department of Biology and Ecology of Fishes at the Thaer-Institute for Life Sciences at Humboldt University of Berlin for the present experiment. For the rearing prior to any experiment, we used a light regime of 12 h light:12 h darkness that resembles the natural surface habitats and maintained water temperature at 26°C. Prior to experiments, test fish were taken from their stock tanks (80-L) and transferred into 54-L tanks in groups of 20 individuals with equal sex ratio. Those 54-L tanks were covered with black plastic foil and could be run either with a 12:12 L:D light regime (6,000 vs. 0 lux) or in total darkness (0 lux; in case fish were later tested in darkness, see below). Fish were fed twice daily *ad libitum* with TetraMin flake food and live Chironomid larvae. Please note that all our test subjects have been raised under normal 12 h light:12 h dark conditions in the lab, and cave mollies thus might not show exactly the same behaviors compared to their wild counterparts that spend their whole lives in darkness. This was necessary to facilitate maintenance work and to ensure that all fish tested in darkness experienced the same treatment since it is impossible to raise surface mollies in darkness without high mortality rates (Riesch et al., 2011). Nevertheless, we acclimated those fish tested in darkness (surface and cave mollies) to complete dark conditions for 1 week prior to our tests. Such an approach follows previous protocols for that species (Plath et al., 2003, 2004).

The BR: RoboFish

Our RoboFish system consists of a glass tank (88 cm × 88 cm) that is filled to a level of 15 cm with aged tap water. The tank is placed on an aluminum rack at about 1.40 m above ground (**Figure 2A**). The two-wheeled differential drive robot moves below the tank on a transparent platform (**Figure 2B**). It carries a neodymium magnet directed to the bottom side of the tank. A three-dimensional (3D)-printed fish replica (**Figure 2C**) is attached to a magnetic base, which aligns with the robot. Hence, the replica can be moved directly by the robot (**Figure 2B**).

FIGURE 1 | Map of our study system. Both tested molly populations originate from the South of Mexico near the city of Tapijulapa, federal state of Tabasco. Here, ancestral forms of *Poecilia mexicana* colonized both surface (b, surface-dwelling molly) as well as cave (a, cave molly) habitats.
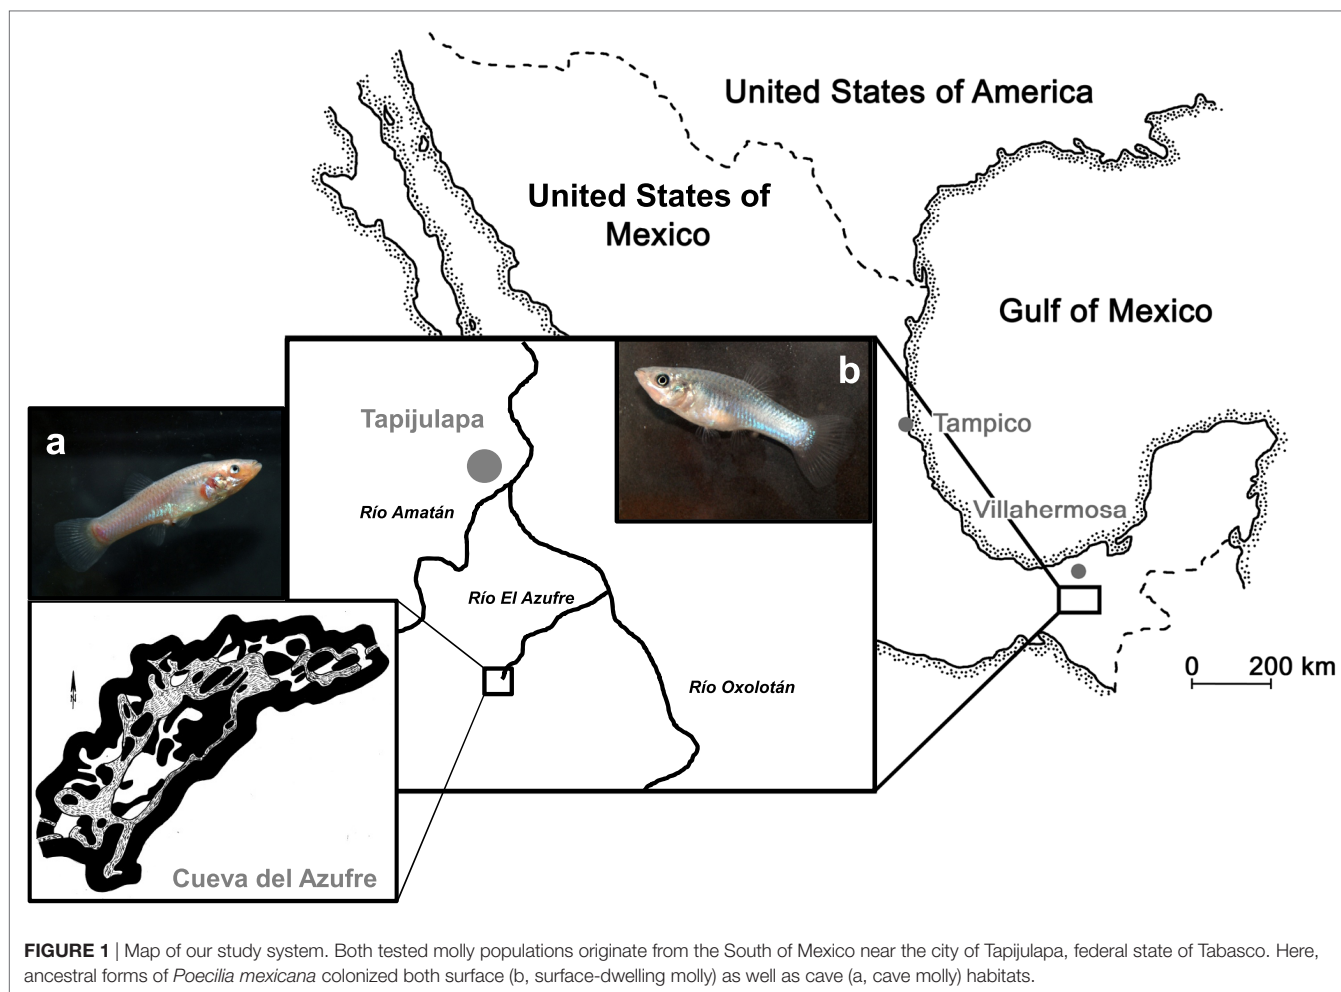
On the ground, a camera is facing upward to track the robot. A second camera (IR-sensitive Bosch Dinion 1080p) is fixed above the tank to track both live fish and replica. The entire system is enclosed in a black, opaque canvas to minimize exposure to external disturbances. For trials in light, the tank was illuminated from above with artificial LED lights reproducing the daylight spectrum (2,000 lux). For trials in darkness, we used four IR spots to light the tank, which cannot be perceived by the fish (Körner et al., 2006) but allows our above-tank camera to record. Two personal computers are used for system operation: one PC tracks (bottom camera) and steers the robot *via* Wi-Fi, whereas a second PC records the video feed of the top camera. The RoboFish moves on a predefined trajectory through the test tank (also called "open-loop" steering). The trajectory used in all described experiments is given in **Figure 3**; RoboFish swims on a continuous zigzag path through the tank. For more detailed information on RoboFish operation modes and construction, see the study by Landgraf et al. (2016).

Experimental Setup: Social Interactions under Two Different Light Conditions

To investigate how surface- and cave-dwelling populations of the Atlantic molly differ in their social behavior in both light

and dark conditions, we observed the interactions of live fish with either RoboFish ($n = 3$ live fish tested for each population and light regime) or with another live conspecific ($n = 3$ pairs of live fish for each population and light regime). The 3D-printed fish replica was modified to match the appearance of *P. mexicana* (**Figure 2C**). The size of the replica (SL: 35 mm) was derived from the mean standard length of all test fish (ranging between 28.77 and 49.03 mm). The replica was situated 0.5 cm below the water surface in accordance with the close-to-surface swimming behavior observed for both populations in the wild (Jourdan et al., 2014). We programmed RoboFish to an average speed of 10 cm/s (maximum speed of 27 cm/s). This is comparable to average speeds obtained for live fish in pilot experiments. The RoboFish swimming sequence was initiated immediately upon transferring the fish into the arena. We started to score social interactions for 3 min after both subjects were first within a range of four body lengths, a distance often assumed to indicate social interactions in poeciliid and other fishes (Croft et al., 2008). Similarly, during trials with conspecifics, two fish were transferred into the arena simultaneously and scoring for 3 min started when fish were moving and after being within a range of 12 cm (ca. 4 body lengths). The fish's movements were tracked using EthoVision™ XT10.1 software
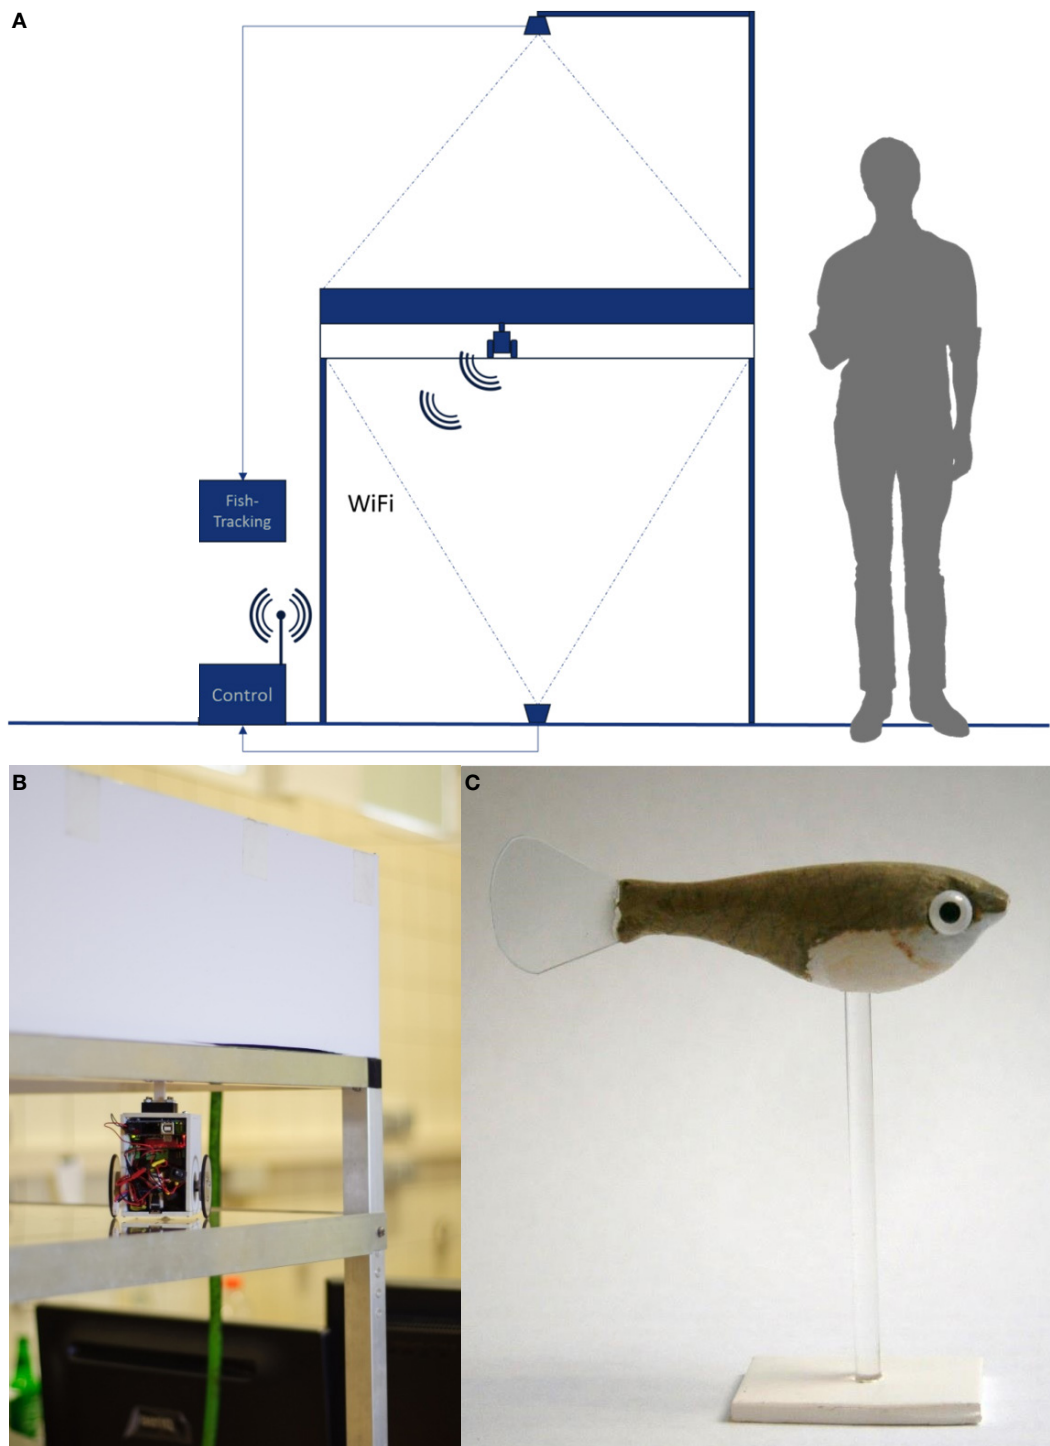
FIGURE 2 | The RoboFish system. **(A)** Experimental setup showing the test tank and bottom as well as top view cameras. The robot is running on a transparent second level below the test tank and is connected via Wi-Fi to the controlling computers. **(B)** Robot close-up below the test tank. **(C)** A molly like replica equipped with glass eyes.

(Noldus Information Technology), and the obtained XY position data were analyzed using customized Python scripts (Python Software Foundation).

Statistical Analysis

Our first aim is to establish whether our focal live fish were socially attracted to their respective companions (either another live fish

FIGURE 3 | Example tracks of a surface-dwelling molly with RoboFish (top) and a live companion (bottom). The swam trajectories (left), the interindividual distance during the 180-s observation period (center) as well as its distribution (right) are shown.

FIGURE 4 | Social behavior of surface and cave-dwelling mollies tested with live conspecifics and RoboFish in light **(A)** and darkness **(B)**. Shown are median interindividual distances along with the results of *U*-tests (*P* values above bars) comparing cave- and surface-dwelling mollies in each treatment. Gray bars represent median and range of simulated interindividual distances for each treatment, and asterisks indicate a significant difference between simulated and real data in Wilcoxon's rank tests (*P* < 0.05).

or RoboFish) in light and darkness. To do so, we compared average median distances between both subjects ("interindividual distance"; see **Figure 3**) in our trials to average median distances obtained for simulated random tracks ("null models"). To obtain "null models", we randomly shuffled focal fish's XY positions at each sampled time step (e.g., by randomly changing order of time steps) and afterward calculated distance between focal fish's XY position to that of the companion's XY position for all time steps. Doing so kept all focal fish's positions but links them randomly with those of the companion. Average medians of interindividual distances of real and simulated tracks were then compared *via* Wilcoxon's rank test (one-tailed, null models are assumed to have greater median interindividual distances), separated by species, companion, and light treatment.

Our second aim is to establish whether surface- and cave-dwelling Atlantic mollies differ in their social behavior (e.g., their shoaling tendency in pairs of live fish or their following tendencies toward RoboFish) and whether light conditions differentially affect social behavior of surface- and cave-dwelling fish. Therefore, we compared average medians of interindividual distances between surface- and cave-dwellers using Mann–Whitney U-tests, separated by light treatment (in light or darkness) and social partner treatment (live companion or RoboFish). Please note that our sample sizes are quite small ($N = 3$ per treatment), which is due to the intense tracking efforts under dark conditions and limited numbers of fish available. Thus, non-significant differences can be a result of low statistical power (in our case beta errors of non-significant tests ranged between 0.05 and 0.40). However, in case of non-significant differences, values were always overlapping.

RESULTS

In light, both cave- and surface-dwelling mollies were similarly strongly attracted to live companions and RoboFish. This was evidenced by significantly smaller interindividual distances among subjects in real interactions compared to simulated tracks (e.g., rank tests comparing median interindividual distances in real interactions and simulated "null models" were significant; see **Figure 4A**). Also, there was no significant difference detectable between the interindividual distances obtained from cave- or surface-dwelling fish when tested with a live companion or RoboFish (U-tests non-significant; see **Figure 4A**). We provide example tracks and interindividual distance plots for RoboFish and live–live interactions of a surface molly in light in **Figure 3**.

In darkness, cave mollies were not attracted to either live companions or RoboFish. Consequently, we found no significant difference between real and simulated tracks (rank tests not significant; **Figure 4B**). Interestingly, surface mollies still showed a strong social attraction toward live companions in darkness with significantly smaller interindividual distances compared to simulated random tracks (**Figure 4B**). However, as seen in cave mollies, surface mollies were not attracted to RoboFish (rank test not significant; **Figure 4B**). Thus, despite our low overall sample size, we found significant differences between cave- and surface-dwelling fish in regard to their social behavior with a live companion in darkness (significance U-test; **Figure 4B**) but not Robofish (non-significance U-test; **Figure 4B**).

DISCUSSION

In light, we found both surface- and cave-dwelling mollies to be similarly strongly attracted by a live conspecific, which contradicts the previously proposed reduced sociality of cave mollies (Plath and Schlupp, 2008). As found in tests with live conspecifics, both ecotypes were following closely a moving BR—RoboFish. This shows the utility of BRs for the study of collective behavior especially in poeciliid fishes (Polverino et al., 2013; Landgraf et al., 2016). However, when tested in darkness, both ecotypes did not follow RoboFish, suggesting that our BR was providing sufficient social cues only when visual inspection was possible. Hence, robotically driven replicas as used in our experiments seem to exploit exclusively visual communication channels. Interestingly, cave fish were also no longer attracted by a live conspecific when tested in darkness, whereas surface-dwellers still showed a significant attraction toward live conspecifics. This contrasts our initial prediction that predominately cave fish with their increased non-visual sensing (Parzefall, 2001) should be able to maintain some degree of sociality also in the dark.

Plath and Schlupp (2008) found that cave mollies from two independently colonized caves (including the population from the Cueva del Azufre also used in our experiments) showed reduced shoaling tendencies when either only visual (stimulus group was presented behind a glass barrier) or both visual and non-visual communications (group presented behind a mesh-wired barrier) was allowed. Thus, the authors concluded that "observed reduction in shoaling in the two cave populations represents a parallel evolutionary process" (Plath and Schlupp, 2008). So, why are cave mollies similarly attracted by live conspecifics and RoboFish compared to surface-dwelling mollies when tested in light in our full-contact experiments? The assumed low sociality of cave mollies was based on dichotomous choice tests in light in which cave- and surface-dwellers had to choose among a group of conspecifics or an empty compartment in the test aquarium (Plath and Schlupp, 2008). While this is a classic and commonly used method to establish shoaling tendencies in small fish (Wright and Krause, 2006), we argue that full contact designs as in our study might lead to different results (Ziege et al., 2012). In addition, technological advances make it easier for the experimenter and thus more common to track animals' movements while they interact unconstrained (Herbert-Read et al., 2011, 2012; Katz et al., 2011; Jolles et al., 2017). Future studies should then focus on comparative approaches evaluating strengths and short comings of either method.

While our tests in light provided cave and surface fish with both visual and non-visual cues and each ecotype might have predominately used one or the other to associate with a live or artificial companion, our tests in darkness omitted visual communication. In experiments using mesh-wired barriers in dichotomous choice tests under dark conditions, Plath et al. (2004, 2005) found only cave mollies to be able to exercise mate choice, an ability that was also confirmed in the wild (Bierbach et al., 2013a). This was attributed to cave mollies exhibiting evolutionary acquired enhanced lateral line (Parzefall, 2001; Parzefall et al., 2007) as well as chemical sensing of conspecifics (Rüschenbaum and Schlupp, 2013; Jourdan et al., 2016). Thus, we initially hypothesized that cave mollies, although assumed to

have an inherent weaker social tendency, should show stronger social attraction in darkness compared to surface fish. We found the opposite with cave mollies showing no social attraction but surface-dwellers were still significantly attracted to a live companion. Shoaling is a behavioral adaptation to predation risk (Krause and Ruxton, 2002), which is strongly reduced in the cave habitat. The Cueva del Azufre is free of piscivorous fish as well as birds, and the only predators preying upon cave mollies are giant water bugs of the genera *Belostoma* (Tobler et al., 2007) as well as pisaurid and theraphosid spiders (Horstkotte et al., 2010) and freshwater crabs (Klaus and Plath, 2011). All these species are sit-and-wait predators that prey from the pool edges and thus have only very limited attack ranges. Thus, it is likely that shoaling does not provide mollies with antipredator benefits in the cave, and there is no evolutionary pressure to maintain shoaling behavior by cave mollies in darkness. This view is supported by experiments showing that cave mollies exhibit reduced avoidance when confronted with fish predators (Bierbach et al., 2013b). It is also possible that cave mollies context dependently adjust their shoaling tendencies in darkness but not in light. This seems to be a unique feature of cave fish as surface-dwelling fish, also habituated to darkness for 1 week, still showed significant shoaling tendencies, probably by using non-visual communication channels like lateral line sensing and conspecific chemical cues (see above). As surface fish might experience predation also in darkness (e.g., during night), maintaining shoaling under dark conditions can be still beneficial. As our sample size was small (see methods) and thus statistical evaluation limited, we recommend future studies to focus on shoaling differences of surface- and cave-dwelling mollies using up-to-date full contact designs and position tracking approaches.

As both cave- and surface-dwelling mollies did not show any social attraction toward RoboFish in darkness, we conclude that our replica is providing only sufficient visual cues but lack other non-visual ones that are important to be recognized as a conspecific in darkness. It is known that tail beating of fish replicas can enhance acceptance probably by stimulating the lateral line system (Marras and Porfiri, 2012), and it seems that a pure swimming (even with direction changes as in our zig-zagged trajectories) does not provide enough similar stimulation. In non-visually communicating animals like weak-electric fishes or insects, researchers tried to mimic species-specific cues by either rebuilding electric discharges at the replica (Donati et al., 2016) or by applying conspecific odors to the replica (Halloy et al., 2007; Landgraf et al., 2011). Furthermore, some researchers now focus on the development of replicas that provide multiple cues (Shi

et al., 2013; Phamduy et al., 2014; Donati et al., 2016; Romano et al., 2017). Future research might focus on exploring which non-visual cues are important for poecillid fishes by step-wise equipping replicas with different artificial cues and comparing the response of live fish in light and darkness. In addition, a comparison with other cave fish will be fruitful as well since several cave ecotypes are blind and thus exclusively rely on non-visual cues (Jeffery et al., 2003). Overall, RoboFish (and similar biomimetic systems) can be a strong tool to investigate social behavior of fish in a standardized way.

ETHICAL STATEMENT

AUTHOR CONTRIBUTIONS

DB, JL, AB, KE, LH, CW, NW, TL, PR, and JK designed the study. DB, JL, LA-R, and JK caught the fish. DB, HM, HN, and TL built the robot system. DB, JL, AB, KE, LH, CW, NW, and HN performed the experiments, DB, KE CW, and PR tracked the videos. DB analyzed the data. All authors interpreted the data and approved the submitted manuscript version.

ACKNOWLEDGMENTS

FUNDING

REFERENCES

Bierbach, D., Penshorn, M., Hamfler, S., Herbert, D. B., Appel, J., Meyer, P., et al. (2013a). Gradient evolution of body colouration in surface- and cave-dwelling *Poecilia mexicana* and the role of phenotype-assortative female mate choice. *Biomed Res. Int.* 2013, 15. doi:10.1155/2013/148348

Bierbach, D., Schulte, M., Herrmann, N., Zimmer, C., Arias-Rodriguez, L., Indy, J. R., et al. (2013b). Predator avoidance in extremophile fish. *Life* 3, 161–180. doi:10.3390/life3010161

Burnett, S. (2011). Perceptual worlds and sensory ecology. *Nat. Educ. Knowl.* 3, 75.

Butail, S., Abaid, N., Macrì, S., and Porfiri, M. (2015). "Fish–robot interactions: robot fish in animal behavioral studies," in *Robot Fish*, eds R. Du, Z. Li, K. Youcef-Toumi, and P. Valdivia y Alvarado (Berlin Heidelberg: Springer), 359–377.

Candolin, U. (2003). The use of multiple cues in mate choice. *Biol. Rev.* 78, 575–595. doi:10.1017/s1464793103006158

Chouinard-Thuly, L., Gierszewski, S., Rosenthal, G. G., Reader, S. M., Rieucau, G., Woo, K. L., et al. (2017). Technical and conceptual considerations for using animated stimuli in studies of animal behavior. *Curr. Zool.* 63, 5–19. doi:10.1093/cz/zow104

Croft, D. P., James, R., and Krause, J. (2008). *Exploring Animal Social Networks*. Princeton: Princeton University Press.

Donati, E., Worm, M., Mintchev, S., Wiel, M. V. D., Benelli, G., Emde, G. V. D., et al. (2016). Investigation of collective behaviour and electrocommunication in the weakly electric fish, *Mormyrus rume*, through a biomimetic robotic dummy fish. *Bioinspir. Biomim.* 11, 066009. doi:10.1088/1748-3190/11/6/066009

Eifert, C., Farnworth, M., Schulz-Mirbach, T., Riesch, R., Bierbach, D., Klaus, S., et al. (2014). Brain size variation in extremophile fish: local adaptation versus phenotypic plasticity. *J. Zool.* 295, 143–153. doi:10.1111/jzo.12190

Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., et al. (2007). Social integration of robots into groups of cockroaches to control self-organized choices. *Science* 318, 1155–1158. doi:10.1126/science.1144259

Harcourt, J. L., Ang, T. Z., Sweetman, G., Johnstone, R. A., and Manica, A. (2009). Social feedback and the emergence of leaders and followers. *Curr. Biol.* 19, 248–252. doi:10.1016/j.cub.2008.12.051

Herbert-Read, J. E., Krause, S., Morrell, L. J., Schaerf, T. M., Krause, J., and Ward, A. J. W. (2012). The role of individuality in collective group movement. *Proc. Biol. Sci.* 280, 1752. doi:10.1098/rspb.2012.2564

Herbert-Read, J. E., Perna, A., Mann, R. P., Schaerf, T. M., Sumpter, D. J. T., and Ward, A. J. W. (2011). Inferring the rules of interaction of shoaling fish. *Proc. Natl. Acad. Sci. U.S.A.* 108, 18726–18731. doi:10.1073/pnas.1109355108

Horstkotte, J., Riesch, R., Plath, M., and Jäger, P. (2010). Predation on a cavefish (*Poecilia mexicana*) by three species of spiders in a Mexican sulfur cave. *Bullet. Br. Arachnol. Soc.* 15, 55–58. doi:10.13156/arac.2010.15.2.55

Jeffery, W. R., Strickler, A. G., and Yamamoto, Y. (2003). To see or not to see: evolution of eye degeneration in Mexican blind cavefish. *Integr. Comp. Biol.* 43, 531–541. doi:10.1093/icb/43.4.531

Jolles, J. W., Boogert, N. J., Sridhar, V. H., Couzin, I. D., and Manica, A. (2017). Consistent individual differences drive collective behavior and group functioning of schooling fish. *Curr. Biol.* 27, 2862–2868.e2867. doi:10.1016/j.cub.2017.08.004

Jourdan, J., Bierbach, D., Riesch, R., Schießl, A., Wigh, A., Arias-Rodriguez, L., et al. (2014). Microhabitat use, population densities, and size distributions of sulfur cave-dwelling *Poecilia mexicana*. *PeerJ* 2, e490. doi:10.7717/peerj.490

Jourdan, J., Jordan, M., Zimmer, C., Eifert, C., Arias-Rodriguez, L., and Plath, M. (2016). Response to conspecific chemical cues in surface- and cave-dwelling populations of Atlantic mollies, *Poecilia mexicana*. *Environ. Biol. Fish.* 99, 697–703. doi:10.1007/s10641-016-0510-0

Katz, Y., Tunstrøm, K., Ioannou, C. C., Huepe, C., and Couzin, I. D. (2011). Inferring the structure and dynamics of interactions in schooling fish. *Proc. Natl. Acad. Sci. U.S.A.* 108, 18720–18725. doi:10.1073/pnas.1107583108

Klaus, S., and Plath, M. (2011). Predation on a cave fish by the freshwater crab *Avotrichodactylus bidens* (Bott, 1969) (Brachyura, Trichodactylidae) in a Mexican sulfur cave. *Crustaceana* 84, 411–418. doi:10.1163/001121611X560853

Körner, K. E., Schlupp, I., Plath, M., and Loew, E. R. (2006). Spectral sensitivity of mollies: comparing surface- and cave-dwelling Atlantic mollies, *Poecilia mexicana*. *J. Fish Biol.* 69, 54–65. doi:10.1111/j.1095-8649.2006.01056.x

Krause, J., and Ruxton, G. D. (2002). *Living in Groups*. Oxford: Oxford University Press.

Krause, J., Winfield, A. F. T., and Deneubourg, J.-L. (2011). Interactive robots in experimental biology. *Trends Ecol. Evol.* 26, 369–375. doi:10.1016/j.tree.2011.03.015

Landgraf, T., Bierbach, D., Nguyen, H., Muggelberg, N., Romanczuk, P., and Krause, J. (2016). RoboFish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live *Trinidadian guppies*. *Bioinspir. Biomim.* 11, 015001. doi:10.1088/1748-3190/11/1/015001

Landgraf, T., Rojas, R., Nguyen, H., Kriegel, F., and Stettin, K. (2011). Analysis of the waggle dance motion of honeybees for the design of a biomimetic honeybee robot. *PLoS ONE* 6:e21354. doi:10.1371/journal.pone.0021354

Magurran, A. E. (2005). *Evolutionary Ecology: The Trinidadian Guppy*. Oxford: Oxford University Press.

Marras, S., and Porfiri, M. (2012). Fish and robots swimming together: attraction towards the robot demands biomimetic locomotion. *J. R. Soc. Interface* 9, 1856–1868. doi:10.1098/rsif.2012.0084

Parzefall, J. (1969). Zur vergleichenden Ethologie verschiedener *Mollienesia*-Arten einschließlich einer Höhlenform von *Mollienesia sphenops*. *Behaviour* 33, 1–38. doi:10.1163/156853969X00297

Parzefall, J. (1970). Morphologische Untersuchungen an einer Höhlenform von *Mollienesia sphenops* (Pisces, Poeciliidae). *Zeitschrift für Morphologie der Tiere* 68, 323–342. doi:10.1007/BF00376005

Parzefall, J. (2001). A review of morphological and behavioural changes in the cave molly, *Poecilia mexicana*, from Tabasco, Mexico. *Environ. Biol. Fish.* 62, 263–275. doi:10.1023/A:1011899817764

Parzefall, J., Kraus, C., Tobler, M., and Plath, M. (2007). Photophilic behaviour in surface- and cave-dwelling Atlantic mollies, *Poecilia mexicana* (Poeciliidae). *J. Fish Biol.* 71, 1225–1231. doi:10.1111/j.1095-8649.2007.01581.x

Peters, N., Peters, G., Parzefall, J., and Wilkens, H. (1973). Über degenerative und konstruktive Merkmale bei einer phylogenetisch jungen Höhlenform von *Poecilia sphenops* (Pisces, Poeciliidae). *Internationale Revue der gesamten Hydrobiologie und Hydrographie* 58, 417–436. doi:10.1002/iroh.19730580308

Phamduy, P., Polverino, G., Fuller, R. C., and Porfiri, M. (2014). Fish and robot dancing together: bluefin killifish females respond differently to the court-ship of a robot with varying color morphs. *Bioinspir. Biomim.* 9, 036021. doi:10.1088/1748-3182/9/3/036021

Plath, M., Heubel, K. U., de Leon, F. G., and Schlupp, I. (2005). Cave molly females (*Poecilia mexicana*, Poeciliidae, Teleostei) like well-fed males. *Behav. Ecol. Sociobiol.* 58, 144–151. doi:10.1007/s00265-005-0918-6

Plath, M., Parzefall, J., Körner, K. E., and Schlupp, I. (2004). Sexual selection in darkness? Female mating preferences in surface- and cave-dwelling Atlantic mollies, *Poecilia mexicana* (Poeciliidae, Teleostei). *Behav. Ecol. Sociobiol.* 55, 596–601. doi:10.1007/s00265-003-0750-9

Plath, M., Parzefall, J., and Schlupp, I. (2003). The role of sexual harassment in cave and surface dwelling populations of the Atlantic molly, *Poecilia mexicana* (Poeciliidae, Teleostei). *Behav. Ecol. Sociobiol.* 54, 303–309. doi:10.1007/s00265-003-0625-0

Plath, M., Pfenninger, M., Lerp, H., Riesch, R., Eschenbrenner, C., Slattery, P. A., et al. (2013). Genetic differentiation and selection against migrants in evolutionarily replicated extreme environments. *Evolution* 67, 2647–2661. doi:10.1111/evo.12133

Plath, M., and Schlupp, I. (2008). Parallel evolution leads to reduced shoaling behavior in two cave dwelling populations of Atlantic mollies (*Poecilia mexicana*, Poeciliidae, Teleostei). *Environ. Biol. Fish.* 82, 289–297. doi:10.1007/s10641-007-9291-9

Polverino, G., Phamduy, P., and Porfiri, M. (2013). Fish and robots swimming together in a water tunnel: robot color and tail-beat frequency influence fish behavior. *PLoS ONE* 8:e77589. doi:10.1371/journal.pone.0077589

Riesch, R., Plath, M., and Schlupp, I. (2011). Speciation in caves: experimental evidence that permanent darkness promotes reproductive isolation. *Biol. Lett.* 7, 909–912. doi:10.1098/rsbl.2011.0237

Romano, D., Benelli, G., Donati, E., Remorini, D., Canale, A., and Stefanini, C. (2017). Multiple cues produced by a robotic fish modulate aggressive behaviour in Siamese fighting fishes. *Sci. Rep.* 7, 4667. doi:10.1038/s41598-017-04840-0

Rüschenbaum, S., and Schlupp, I. (2013). Non-visual mate choice ability in a cavefish (*Poecilia mexicana*) is not mechanosensory. *Ethology* 119, 368–376. doi:10.1111/eth.12072

Shi, Q., Ishii, H., Kinoshita, S., Takanishi, A., Okabayashi, S., Iida, N., et al. (2013). Modulation of rat behaviour by using a rat-like robot. *Bioinspir. Biomim.* 8, 046002. doi:10.1088/1748-3182/8/4/046002

Tinbergen, N. (1948). Social releasers and the experimental method required for their study. *Wilson Bull.* 60, 6–51. doi:10.2307/4157642

Tobler, M., Palacios, M., Chapman, L. J., Mitrofanov, I., Bierbach, D., Plath, M., et al. (2011). Evolution in extreme environments: replicated phenotypic differentiation in livebearing fish inhabiting sulfidic springs. *Evolution* 65, 2213–2228. doi:10.1111/j.1558-5646.2011.01298.x

Tobler, M., Schlupp, I., and Plath, M. (2007). Predation of a cave fish (*Poecilia mexicana*, Poeciliidae) by a giant water-bug (*Belostoma*, Belostomatidae) in a Mexican sulphur cave. *Ecol. Entomol.* 32, 492–495. doi:10.1111/j.1365-2311.2007.00892.x

Webb, B. (2000). What does robotics offer animal behaviour? *Anim. Behav.* 60, 545–558. doi:10.1006/anbe.2000.1514

Wright, D., and Krause, J. (2006). Repeated measures of shoaling tendency in zebrafish (*Danio rerio*) and other small teleost fishes. *Nat. Protoc.* 1, 1828–1831. doi:10.1038/nprot.2006.287

Ziege, M., Hennige-Schulz, C., Muecksch, F., Bierbach, D., Tiedemann, R., Streit, B., et al. (2012). A comparison of two methods to assess audience-induced changes in male mate choice. *Curr. Zool.* 58, 84–94. doi:10.1093/czoolo/58.1.84

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

A Two Teraflop Swarm

Simon Jones[1,2,3]*, Matthew Studley[2,3], Sabine Hauert[1,3] and Alan Frank Thomas Winfield[2,3]

[1]University of Bristol, Bristol, United Kingdom, [2]University of the West of England, Bristol, United Kingdom, [3]Bristol Robotics Laboratory, University of the West of England, Bristol, United Kingdom

We introduce the Xpuck swarm, a research platform with an aggregate raw processing power in excess of two teraflops. The swarm uses 16 e-puck robots augmented with custom hardware that uses the substantial CPU and GPU processing power available from modern mobile system-on-chip devices. The augmented robots, called Xpucks, have at least an order of magnitude greater performance than previous swarm robotics platforms. The platform enables new experiments that require high individual robot computation and multiple robots. Uses include online evolution or learning of swarm controllers, simulation for answering *what-if* questions about possible actions, distributed super-computing for mobile platforms, and real-world applications of swarm robotics that requires image processing, or SLAM. The teraflop swarm could also be used to explore swarming in nature by providing platforms with similar computational power as simple insects. We demonstrate the computational capability of the swarm by implementing a fast physics-based robot simulator and using this within a distributed island model evolutionary system, all hosted on the Xpucks.

Keywords: swarm robotics, robot hardware, simulation, evolutionary robotics, behavior trees, distributed evolutionary algorithm, GPGPU, embodied reality modelling

1. INTRODUCTION

The Xpuck swarm is a new research platform with an aggregate raw processing power in excess of two teraflops, which enables new experiments that require high-individual robot computation and large numbers of robots. There are several research areas that particularly motivate the design.

Swarm robotics (Sahin, 2005) originally takes inspiration from collective phenomena in nature, including social insects, flocks of birds, and schools of fish to create collective behaviors that emerge from local interactions between robots and their environment. These swarms have the potential to be inherently robust, decentralized, and scalable. A fundamental problem of the field is the automatic design of controllers for robot swarms such that a desired collective behavior emerges (Francesca and Birattari, 2016). One common and successful approach is the use of evolutionary techniques to discover suitable controller solutions in simulated environments and the transfer of these controllers to real robots. However, this often results in lower performance due to the *reality gap* (Jakobi et al., 1995). Embodied evolutionary swarm robotics moves evolution into the swarm and directly tests controllers, avoiding the reality gap and making the swarm scalable and adaptive to the environment (Watson et al., 2002). Usually, the low-processing power of the individual robots precludes using simulation within the robots as a means of accelerating the evolutionary process. Moving computational power into the swarm would allow us to combine these approaches, the speed of evolution within simulated environments together with the adaptability of continuous reality testing.

Giving a robot the ability to answer *what-if* questions could allow a robot to evaluate courses of action or strategies in the safety of simulation, rather than in the real world where they may have

TABLE 1 | Current and potential swarm platforms.

Product	SoC	GFLOPS (fp32)	RAM (bytes)	Price (£)
Robot platforms				
Kilobot	Atmel atmega328p	0.0008[a]	2 K	15
e-puck	dsPIC	0.0015[a]	8 K	650
r-one	TI Stellaris LM3S8962	0.005	64 K	165
Linux Extension Board	Atmel AT91SAM9260	0.02[a]	64 M	80[i]
Swarmbots	Intel Xscale	0.04[a]	64 M	Not known
GCTronic Gumstick	TI AM3703	1.2	512 M	600[i]
Khepera IV	TI OMAP3730	1.2	512 M	2,000
Pi-puck	Broadcom BCM2835	1.4[b]	512 M	110[i]
Pheeno	Broadcom BCM2836	7.2[c]	1 G	205
Xpuck	Samsung Exynos 5 Octa (5422)	36 + 122[d]	2 G	135
Compute modules				
Hardkernel XU4	Samsung Exynos 5 Octa (5422)	36 + 122[d]	2 G	70
Samsung Artik 1020	Exynos 5 Octa[e]	36 + 122[d]	2 G	98
Wandboard IMX6Q	NXP i.MX6 Quad	25[f]	2 G	120
Intrinsyc Open-Q820SOM	Qualcomm Snapdragon 820	250[g]	3 G	250
Nvidia Jetson TX1	Nvidia Tegra 210	512[h]	2 G	290

[a]Integer only, assumes 10 integer instructions per floating point operation.
[b]VMLA × 0.7 GHz. VideoCore IV GPU has no OpenCL support.
[c]VMLA × 4 × 0.9 GHz. VideoCore IV GPU has no OpenCL support.
[d]CPUs A7 1.4 GHz, A15 0.8 GHz + ARM Mali-T628MP6 GPU, 4 vector multiplies, 4 vector adds, 1 scalar multiply, 1 scalar add, 1 dot product per cycle, 6 cores, each with 2 arithmetic pipelines at 600 MHz. OpenCL 1.2 full profile.
[e]Assumption. The product literature does not state the SoC but Samsung only used the Mali-T628MP6 in the Exynos 5 Octa family.
[f]Vivante GC2000 GPU only, 4 vector multiplies, 4 vector adds, 4 cores at 794 MHz, OpenCL 1.1 embedded profile.
[g]Very little open information, https://en.wikipedia.org/wiki/Adreno states 498.5 at 624 MHz but assumed to be fp16 rather than fp32. OpenCL 2.0.
[h]According to AnandTech, Ho and Smith (2015).
[i]In addition to e-puck cost.

potentially catastrophic consequences. The utility of this ability depends on the speed of simulation; clearly the higher the speed, the more possibilities can be tested. One use of internal *what-if* modeling is the "ethical" robot of Winfield et al. (2014), which uses simulation to allow a robot to predict the consequences of its actions or inactions on other agents and choose an ethical course of action. Another use of internal reality modeling is to detect faulty or corrupted members of a swarm by noticing deviations from predicted behavior. For safety critical applications, or where the potential consequences of actions are serious, using an unreliable communications link to remote systems would not be possible and the embodiment of the simulation within the robot is essential.

A third intriguing area where increased computational ability could be applied is in much more complex neural net controllers. Although swarm robotics as a field is inspired by social insects and other animals, the robot agents are far simpler than the organisms which inspire their creation. As a crude example, the number of neurons in an ANN controller for a swarm system rarely exceeds a dozen. Neurons in animal brains are considerably more complex and numerous; the nematode worm *C. elegans* has 302, the parasitic wasp *Megaphragma mymaripenne* has 7,400, an ant has 2.5×10^5, and a honey bee has a million (White et al., 1986; Menzel and Giurfa, 2001; Polilov, 2012). The system we describe could simulate several thousand biologically plausible neurons per Xpuck.

These three areas would benefit from greatly increased processing power within the robots of a swarm, enabling either simulation of physical systems or execution of complex controllers. Many other applications of robotics such as SLAM or image processing also require high-processing power. Consumer electronics has

been improving in performance for many years. Moore's Law (Mack, 2011) observes that the number of transistors for a given cost is doubling every 18 months and their power consumption is decreasing in proportion. Over 10 years, we should expect to see a given processing performance become available with one hundredth the power consumption.[1] This makes it now possible to build a high-computing performance swarm running on limited battery power.

In this paper, we describe the design of new swarm robotics platform that makes use of this recently available and cheap high-performance computing capability to augment the widely used e-puck robot, which many labs will already have available. We have designed it to have higher computational capability than any other swarm platforms, see **Table 1**, and to have a battery life at least as good as other solutions, while minimizing costs to allow the building of large swarms. We demonstrate the computational capability of the platform in two ways. First, we evaluate a fiducial tracking image processing application using the e-puck camera that would not be computationally possible on the standard e-puck. Second, and to lay the groundwork for future experiments, we implement a fast parallel physics-based robot simulator running on the GPU of the Xpuck, and use this within a distributed island-model evolutionary system to discover swarm controllers.

2. MATERIALS AND METHODS

In this section, we set out our system requirements. We outline potential computing modules. We characterize the

---

[1]2004 Nvidia 6800 Ultra 40 GFLOPS 110 W, 0.35 GFLOPS/W. 2014 Samsung Exynos 5422 120 GFLOPS 5 W, 24 GFLOPS/W.

power/performance tradeoffs of our chosen compute module and then discuss the design and implementation of the Xpuck hardware and associated system infrastructure to enable running experiments. We then detail the design and implementation of a fast physics-based robot simulator specifically tailored to the Xpuck to enable onboard evolutionary algorithms. We also describe two demonstrations of the Xpuck computational capabilities, a fiducial tracking application that could not be run on a standard e-puck, and an island model evolutionary algorithm running on multiple Xpucks.

To run experiments building on the literature, we decided that, in addition to much higher processing power, the Xpuck must meet or exceed the capabilities provided by the existing e-puck robots with additional processing boards. The e-puck is a two-wheel stepper motor-driven robot. Its sensors comprise a ring of IR proximity sensors around its periphery, a three-axis accelerometer, three microphones, and a VGA video camera. As with the Linux Extension Board (LEB), introduced by Liu and Winfield (2011), we require a battery life of at least 1.5 h and full access to the e-puck's IR proximity and accelerometer sensors, and control of the stepper motors and LEDs. In addition, we require that the VGA camera can stream full frame at >10 fps. The Xpuck must run a full standard Linux, able to support ROS (Quigley et al., 2009). It must have WiFi connectivity. GPGPU capabilities must be made available through a standard API such as OpenCL or CUDA (Nvidia, 2007; Khronos OpenCL Working Group, 2010). We also want multicolor LED signaling capability for future visual communication experiments (Floreano et al., 2007; Mitri et al., 2009). Since many labs already have multiple e-puck robots, we wished to minimize the additional cost of the Xpuck to facilitate the construction of relatively large swarms of robots. With this in mind, we chose a target budget per Xpuck of £150.

Given the requirements, **Table 1** sets out some of the current swarm platforms and potential modules that could be used to enhance the e-puck. There are a number of interesting devices, but unfortunately there are very few that are commercially available at a budget suitable to satisfy the cost requirement of £150. Within these cost constraints, of the two Samsung Exynos 5 Octa-based devices, the Hardkernel XU4 and the Samsung Artik 1020, only the XU4 was more widely available at the time of design. The Artik module became generally available in early 2017 and would be interesting for future work because of its small form-factor. There are other small form-factor low-cost modules such as the Raspberry Pi Zero, as used in the Pi-puck (Millard et al., 2017), but none that provide standard API access to GPGPU capability. For these reasons, we chose to base the Xpuck on the Hardkernel Odroid XU4 single board computer.

2.1. High-Performance Computing

The Hardkernel Odroid XU4 is a small single board computer based around the Samsung Exynos 5422 SoC. It has 2 GB of RAM, mass storage on microSD card, ethernet and USB interfaces, and connectors exposing many GPIO pins with multiple functions.

The SoC contains eight ARM CPU cores in a big.LITTLE[2] formation, i.e., two clusters, one of four small low power A7 cores, and one of four high-performance A15 cores. The system concept

---

[2]https://developer.arm.com/technologies/big-little.

TABLE 2 | Hardkernel Odroid XU4 specifications.

Spec	Details
SoC	Samsung Exynos 5 Octa (5422)
CPU organization	big.LITTLE 4 + 4
CPU big	4× ARM Cortex A15 2 GHz 4 × 32 K L1I,
	4 × 32 K L1D, shared 2 M L2 25.6GFLOPS[a]
CPU little	4× ARM Cortex A7 1.4 GHz 4 × 32 K L1I,
	4 × 32 K L1D, shared 512 K L2 11.2GFLOPS[b]
GPU	ARM Mali T628MP6 600 MHz 122GFLOPS[c]
Memory	2 GB LPDDR3 933 MHz PoP
Memory bandwidth	14.9 GB/s
Idle power	2 W
Maximum power	21 W

[a]4-wide SP NEONv2 FMA × 4 × 800 MHz.
[b]VMLA × 4 × 1.4 GHz.
[c]4 vector multiply, 4 vector add, 1 scalar multiply, 1 scalar add, 1 dot product per cycle × 2 pipelines × 6 cores × 600 MHz.

envisages the small A7 cores being used for regular but undemanding housekeeping tasks, and the higher performing A15 cores being used when the computational requirements exceed that of the A7 cores, at the expense of greater power consumption. It also contains an ARM Mali T628-MP6 GPU, which supports OpenCL 1.2 Main Profile, allowing the relatively easy use of the GPU for GPGPU computation. Some important specifications are detailed in **Table 2**.

The Linux kernel supplied by Hardkernel supports full Heterogeneous MultiProcessor (HMP) scheduling across all eight cores, with the frequencies of the two clusters being varied according to the current process mix and load, the specified minimum and maximum frequencies for each cluster, and the kernel *governor* policy.[3] It was evident from manually changing the CPU frequencies during initial investigation that there was little subjective performance boost from using the highest frequencies, but a large increase in power consumption.

2.1.1. Operating Point Tuning

Computational efficiency is an important metric, directly affecting the battery life. Initial tests showed that setting the maximum frequencies to the highest allowed by the hardware (A15—2 GHz, A7—1.4 GHz) and running a computationally heavy load caused the power consumption to exceed 15 W. To characterize the system and find an efficient operating point, we chose to perform benchmarking with a large single precision matrix multiplication using the standard BLAS API function SGEMM. This computes $C = \alpha AB + \beta C$, which performs $2N^2(N+1)$ operations for an $N \times N$ matrix. Good performance requires both high real floating point performance and good memory bandwidth. The OpenBLAS libraries (Xianyi et al., 2012) provide optimized routines capable of running on multiprocessor systems and can utilize all available processor cores. ARM provides useful application notes on implementing an efficient single precision GEMM on the GPU (Gronqvist and Lokhmotov, 2014).

Power consumption was measured for the XU4 board as a whole, using an INA231 with a 20-mΩ shunt resistor in series with the 5-V supply. A cooling fan attached to the SoC was run continuously from a separate power supply to prevent the fan

---

[3]Essentially, how fast clock frequency will be varied to meet changing CPU load.

control thermal regulation from affecting the power readings. Clock frequency for the A7 and A15 clusters of the Exynos 5422 were varied in 200 MHz steps from 200 MHz to 1.4 GHz for the A7, and from 200 MHz to 2 GHz for the A15 clusters, respectively. At each step, a 1,024 by 1,024 SGEMM was performed continuously and timed for at least 5 s while the power values were measured to give Floating Point Operations per second (FLOPS) and FLOPS/W. All points in the array were successfully measured except for the highest frequency in both clusters; 1.4 GHz for A7 and 2 GHz for A15, which caused the SoC temperature to exceed 95°C during the 5-s window, even with the cooling fan running, resulting in the automatic clock throttling of the system to prevent physical damage.

The results confirm that increasing CPU clock frequencies, particularly of the A15 cluster, produced little performance gain but much higher power consumption. **Figure 1** shows that the most efficient operating point of 1.95 GFLOPS/W and 9.1 GFLOPS occurs at the maximum A7 cluster frequency of 1.4 GHz, and the relatively low A15 cluster frequency of 800 MHz. Increasing the A15 frequency to the maximum achievable of 1.8 GHz results in a 6% increase in performance to 9.7 GFLOPS but at the cost of 40% drop in efficiency to 1.21 GFLOPS/W. Because of this dramatic drop in efficiency, we fix the maximum A15 frequency to 800 MHz.

As with the CPU measurement, GPU power consumption was measured for the system as a whole, in the same way. The clock frequency of the GPU was set to each of the allowed frequencies of 177, 266, 350, 420, 480, 543, and 600 MHz and an OpenCL kernel implementing a cache efficient SGEMM was repeatedly run on both the OpenCL devices. **Figure 1** shows that efficiency only declines slightly from the peak at around 480 MHz to 2.24 GFLOPS/W and 17.7 GFLOPS at the maximum 600 MHz. For this reason, we left the maximum allowed frequency of the GPU unchanged.

Note that the GFLOPS figures in these tests are much lower than the theoretical peak values in **Table 2** because the SGEMM task is mostly memory bound.

2.2. Interface Board

An interface board was created to provide power to the XU4 single board computer, interface between the XU4 and the e-puck, and provide new multicolor LED signaling. The overall structure is shown in **Figure 2**.

There are three interfaces to the e-puck, all exposed through the expansion connectors; a slow I^2C bus that is used for controlling the VGA camera, a fast SPI bus that is used for exchanging data packets between the XU4 and the e-puck, over which sense and control information flow, and a parallel digital interface to the VGA camera. In each case, the interfaces have 3.3v logic levels.

The XU4 board has a 30 pin expansion connector that exposes a reasonable number of the GPIO pins of the Exynos 5422 SoC, some of which can be configured to be I^2C and SPI interfaces. The XU4 interface logic levels are 1.8 V. A camera interface was not available, and initial investigation showed that it would not be possible to use pure GPIO pins as a parallel data input from the camera due to the high required data rate. We decided to use a USB interface to acquire camera data.

We intend to use visual signaling as a means of communication within swarms. For this purpose, we included a ring of fifteen Neopixels around the edge of the interface board. Neopixels are relatively recently available digital multicolor RGB LEDs which are controlled with a serial bitstream. They can be daisy chained in very large numbers and each primary color is controllable to 256 levels.

2.2.1. Power Supply

The XU4 requires a 5-V power supply. To design the power supply, the following constraints are assumed:

- The XU4 and supporting electronics will be powered from their own battery, separate from the e-puck battery.
- The average power consumption will be 5 W.
- The peak power consumption will be 10 W.

It is immediately clear that the e-puck battery, a single-cell Li-ion type with a capacity of about 1,600 mAh, would not be able to power the XU4 as well. At a cell voltage of 3.7 V, converter efficiency of 85% and a nominal power consumption of 5 W, battery life would be at best $\frac{3.7 \times 1.6 \times 0.85}{5} = 1$ hour, not counting the requirements of the e-puck itself. These estimates are based on battery characteristics in ideal conditions and real world values will be lower. Hence, they need for a second battery. To get a 1.5-h endurance, we assume a conservative 50% margin to account for real-world behavior, giving the requirement of $\frac{1.5 \times 1.5 \times 5}{3.7 \times 0.85} = 3.6\ Ah$.

Mobile devices are generally designed to work within a power envelope of around 5 W or the case becomes too hot to hold comfortably, see, for example, Gurrum et al. (2012). We assume that with attention to power usage, it will be possible to keep the average power at this level.

The third constraint was motivated by a survey of the readily available switch-mode power supply solutions for stepping up from 3.7 V single-cell lithium to the required 5 V. Devices tended to fall into two types—boost converters that were capable of high currents (>2 A) but with low efficiencies and large-sized inductors due to low-operating frequencies, or devices designed for mobile devices which include battery protection and have small sized inductors due to their high efficiency and operating frequency. Of the latter class, the highest output current was 2 A, with future higher current devices planned but not yet available. Measurements of the XU4 showed an idle current of 400 mA but very high current spikes, exceeding 3 A during booting. To meet the third constraint and enable the use of a high efficiency converter, the kernel was modified to boot using a low clock frequency, reducing boot current to below 1.5 A.

The power supply regulator chosen was the Texas Instruments TPS61232. It is designed for single-cell Li-ion batteries, has a very high efficiency of over 90%, a high switching frequency of 2 MHz resulting in a physically small inductor, and has battery protection with undervoltage lockout.

One aspect of the power supply design that is not immediately obvious is that the battery current is quite high, reaching 4 A as the cutoff discharge limit of 2.5 V is reached. This seriously constrains switching the input power. In fact, physically small

FIGURE 1 | Performance, power consumption, and efficiency of the CPUs and GPU while continuously running a 1,024 × 1,024 single precision matrix multiplication. Highest efficiency for the CPU clusters is with the maximum A7 frequency of 1.4 GHz but a relatively low A15 frequency of 800 MHz. The GPU efficiency stays relatively flat above 480 MHz.
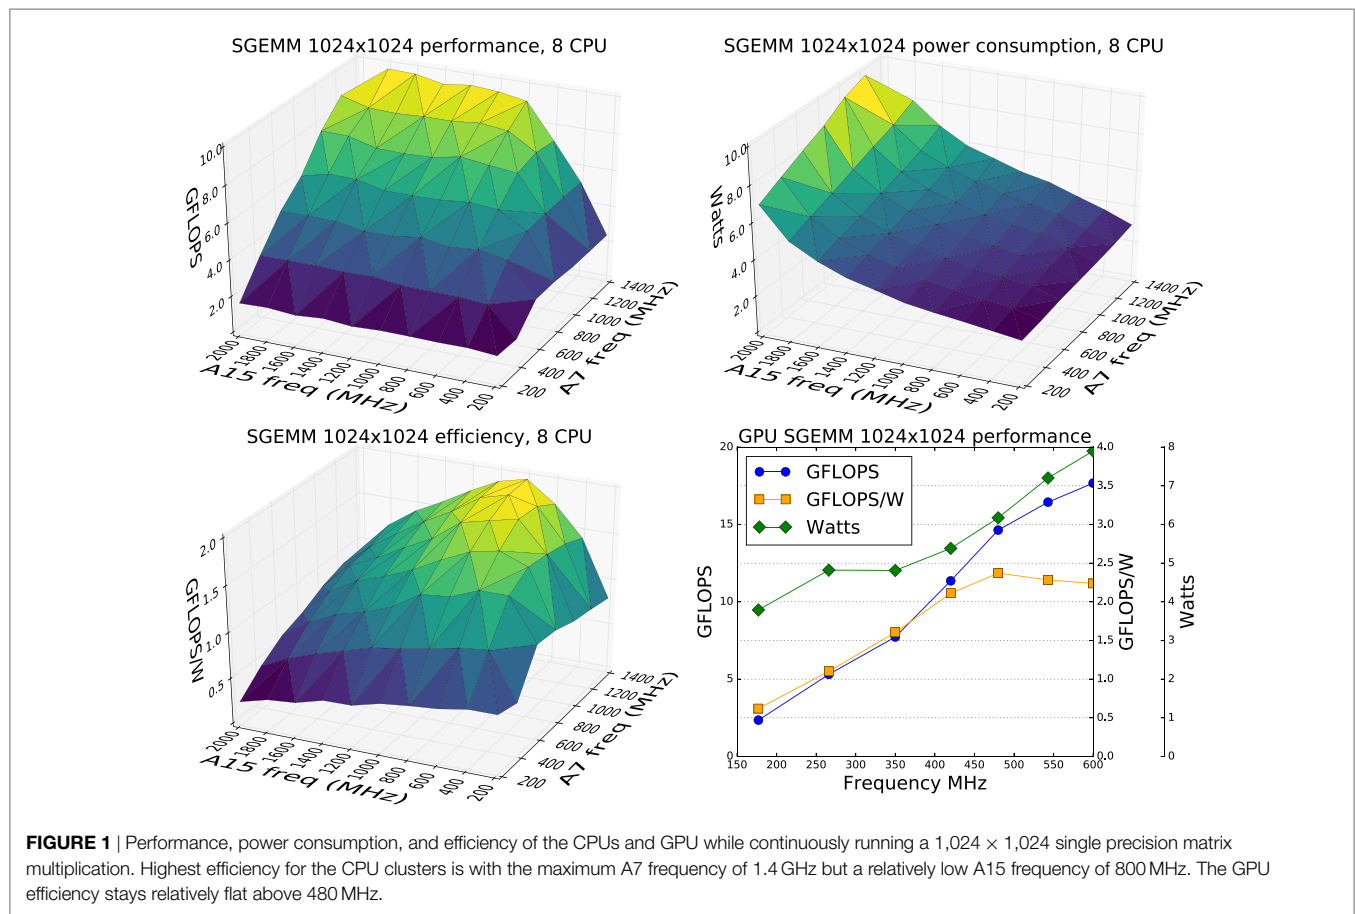
switches capable of handling this amount of current are not readily available. For this reason, and to integrate with the e-puck, two Diodes Incorporated AP2401 high side switches were used in parallel to give electronic switching, allowing the use of the e-puck power signal to turn on the XU4 supply. The high current also necessitates careful attention to the resistance budget and undervoltage lockout settings.

To monitor battery state and energy, we use two Texas Instruments INA231 power monitoring chips, sensing across 20-mΩ resistors on the battery and XU4 side of the switching regulator. These devices perform automatic current and voltage sensing, averaging and power calculation, and are accessible over an I^2C bus. The Hardkernel modified Linux kernel also targets the older Odroid XU3 board, which included the same power monitor chips, so the driver infrastructure is already present to access them.

We used branded Panasonic NCR18650B batteries, rated at 3,400 mAh, and achieved a battery life of close to 3 h while running a ROS graph with nodes retrieving camera data at 640 × 480 pixels 15 Hz, performing simple blob detection, exchanging control packets at 200 Hz with the e-puck dsPIC and conditioning the returned sensor data, and running a simple swarm robot controller. All the LEDs were lit at 50% brightness and varying color, and telemetry was streamed over WiFi at an average bandwidth of 10 kB/s. **Figure 3** shows the discharge curve. Power is relatively constant throughout at about 3.3 W except at the end, where it

drops slightly. This is due to the Neopixel LEDs being supplied directly from the battery. As the voltage drops below about 3.1 V, the blue LEDs stop working, reducing the power consumption.

2.2.2. Camera Interface

The e-puck VGA camera is a Pixelplus PO3030K or PO6030K, depending on the e-puck serial number. Both types have the same electrical interface, although the register interface is slightly different. It is a 640 × 480, 30 fps CMOS sensor, controlled by I^2C, and supplies video on an eight bit parallel bus with some additional lines for H and V sync. By default, the camera provides 640 × 480 data within an 800 × 500 window in CrYCbY format. Each pixel is 16 bits and takes two clocks. The maximum clock frequency of 27 MHz gives 30 fps, with a peak bandwidth of 27 MB/s, sustained 18.4 MB/s. At our minimum desired frame rate of 10 Hz, the clock would be 9 MHz.

We considered a number of possible solutions to the problem of getting the VGA camera data into the XU4, initially focusing on implementing a USB Video Class device, which would then be simply available under the standard Linux webcam driver but available devices were relatively expensive (e.g., XMOS XS1-U8A-64 £18, Cypress Semiconductor CYUSB3014 £35, UVC app notes available for both). In the end, we settled on a more flexible approach, using the widely available and cheap FTDI FT2232 USB interface chip, together with a low power and small FPGA from Lattice.
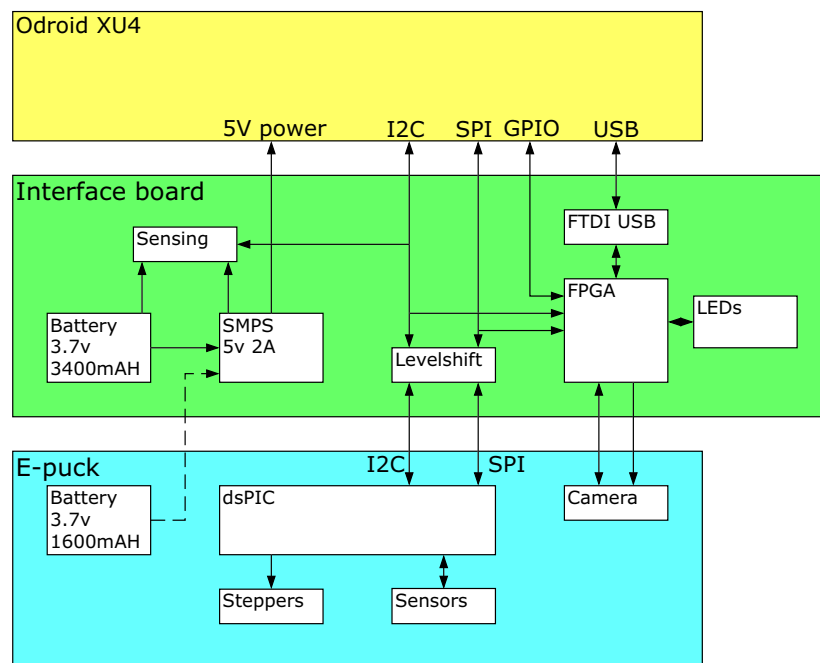
FIGURE 2 | Block diagram showing the functionality of the interface board. The yellow box at the top is the XU4 single board computer, communicating over I²C, SPI, and USB interfaces with the interface board in green. This performs voltage level shifting, provides a USB interface to the e-puck camera, and supplies 5v power to the XU4. The e-puck in blue acts as a slave device to the XU4, running low-level proximity sensing and stepper motor control.

FIGURE 3 | Battery life of close to 3 h while running a ROS graph with nodes retrieving camera data at 640 × 480 pixels 15 Hz, performing simple blob detection, exchanging control packets at 200 Hz with the e-puck dsPIC, and running a basic behavior tree interpreter. All the Neopixel LEDs were lit at 50% brightness and varying color, and telemetry was streamed over WiFi at an average bandwidth of 10 kB/s. The fall-off in power consumption at the 2.5-h point is due to the battery voltage falling below the threshold voltage of the blue LEDs within the Neopixels.

We wanted a low-cost solution; the FT2232H is around £5, and provides a USB2.0 High Speed interface to various other protocols such as synchronous parallel, high speed serial, JTAG, etc. It is not programmable though, and cannot enumerate as a standard UVC device. The FT2232H provides a bulk transfer mode endpoint.

This is not ideal for video, since it provides no latency guarantees, unlike isosynchronous mode, but since we control the whole system, we can ensure that there will be no other devices on the USB bus that could use transfer slots.

Although the FT2232H provides a synchronous parallel interface, it is not directly compatible with the camera. The FT2232H has a small amount of buffering, and uses handshaking to provide backpressure to the incoming data stream if it cannot accept new data, whereas the camera has no storage and simply streams data at the clock rate during the active 640 pixels of each line. To provide buffering and handle interfacing, we chose to use the Lattice Semiconductor iCE40HX1K FPGA. This low-cost device, less than £4 in a TQ144 package, has 96 programmable IO pins in four banks each of which that can run with 1.8, 2.5, or 3.3 V IO standards. It has 64 kB of RAM, sufficient to buffer 6.4 lines of video, or 1.3 ms at our minimum desired frame rate. We assume that the Linux USB driver at the XU4 end can handle all incoming USB data provided there is an available buffer for the data, meaning that the combined maximum latency of the user application and kernel driver must not exceed 1.3 ms to avoid underruns. Given reported sustained data rates of 25 MB/s for the FT2232H, this seems plausible, although should this not prove possible, we had the fallback position of being able to lower the camera clock frequency to a sustainable level.

The decision to use an FPGA with the large number of IOs capable of different voltage standards gave greater design freedom. There is no need for any other glue logic, and it is possible to design defensively, with a number of alternative solutions to each interface problem. It also makes possible the later addition of other peripherals. For this reason, sixteen uncommitted FPGA
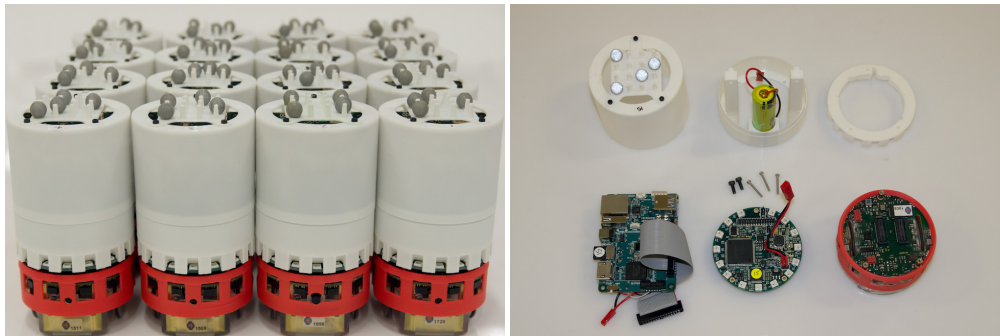
FIGURE 4 | Left: 16 assembled Xpucks. Right: major components, left to right, top to bottom. Outer 3D printed shell showing Vicon tracking reflectors in unique pattern on top. Support chassis, which holds the XU4 single board computer and the LiION battery. Spacer ring, locating the chassis above the PCB and reflecting the upward facing LEDs outwards. XU4 computer, with leads for power and data. Interface PCB. Base e-puck, with red 3D printed skirt.

pins were brought out to an auxiliary connector. Lattice semiconductor provides an evaluation kit, the iCEstick, broadly similar to the proposed subsystem, allowing early development before the completion of the final PCBs.

The final system proved capable of reliably streaming camera data at 15 fps, or 9.2 MB/s, with a camera clock of 12 MHz.

2.2.3. I²C and SPI Communications, Neopixel LEDs

All the e-puck sense and control data, except for the camera, flow over the SPI interface. It is used to control the e-puck motors and LEDs, the Neopixel LEDs on the interface board, and to read from the accelerometers and IR proximity sensors on the e-puck. The I²C bus is only used to set the parameters of the VGA camera.

As with the LEB, the XU4 board acts as the SPI master, providing the clock and enable signals, and the dsPIC of the e-puck the slave. SPI communication is formed of 16-bit packets. Both the master and slave have a 16-bit shift register and communication is full duplex. The master loads data into its register and signals the start of communications, followed by 16 clocks, each shifting one bit of the shift register out from the master and into the slave. Simultaneously, the slave data are shifted into the master. Between each 16 bit packet, communication pauses for long enough for the master and slave to process the received packet and prepare the next outgoing packet. This is handled in hardware with DMA at the XU4 end, but the dsPIC has no DMA and uses an interrupt routine to perform this. We used a value of 6.4 μs to ensure sufficient processing time.

The SPI signals were routed to the FPGA and the board design allows for them to be routed through it. This enables two things: first, the FPGA can watch the data from the XU4 and use fields within that to control its own peripherals, currently the Neopixel LEDs, second, it allows the insertion of data into, or the modification of the return messages from the e-puck.

The FPGA contains additional logic to interpret fields within the SPI packet for controlling the Neopixel LEDs. These data are stored in a buffer within the FPGA and used to generate the appropriately formatted serial stream to the LEDs.

2.3. Physical Design

The interface board is 70 mm in diameter, the same as an e-puck. It sits on top of the base e-puck. Above this, the XU4 board is held

vertically within a 75-mm diameter cylindrical 3D printed shell, which also holds the battery. Flying leads from the XU4 for the GPIO parallel and the USB interfaces, and for the power supply, connect to the interface board. **Figure 4** shows 16 completed Xpucks, and the major components of the assembly. **Figure 5** shows details of a populated interface board.

2.4. Software and Infrastructure

The swarm operates within an infrastructure that provides tracking, virtual sensing, and message services. To facilitate this, the Xpucks run a full featured distribution of Linux and ROS, the Robot Operating System (Quigley et al., 2009). This gives access to much existing work: standard libraries, toolchains, and already existing robot software. Given the close dependence of ROS on Ubuntu we chose to use Ubuntu 14.04.4 LTS, running ROS Indigo.

2.4.1. Real-time Kernel

The standard Linux kernel is not hard real-time, i.e., it does not offer bounded guarantees of maximum latency in response to events. One of the tasks that are running on the XU4 that requires real-time performance is the low-level control loop comprising the SPI data message exchange with the e-puck. The maximum speed of the e-puck is about 130 mm/s. A distance of 5-mm corresponds to about 40 ms. It would be desirable to have a control loop with a period several times faster than that, one commonly used in e-puck experiments is 100 Hz, or $t_{control} = 10$ ms. The minimum time for the control loop to respond to a proximity sensor is two SPI message lengths, so to achieve a 10-ms control period, we need an SPI message period $t_{period} < 5$ ms. Assuming a 5-MHz SPI clock with a message comprising 32 16-bit packets and a 6.4-μs interpacket gap, the total time per message is $t_{message} = 307$ μs. This gives a budget of $t_{period} - t_{message} = 4.7$ ms for processing and latency. Measurements using cyclictest[4] over 500,000 loops of 1 ms, or about 8 min, with the *Server* preemption policy kernel while running SPI message exchange at 200 Hz showed figures of 13.9 ms, and even when running the *Low-Latency Desktop* preemption policy this was above 3.5 ms. This leaves little margin for processing.

We used the PREEMPT-RT patch (Rostedt and Hart, 2007), which modifies the kernel to turn it into a real-time operating

---

[4]https://wiki.linuxfoundation.org/realtime/documentation/howto/tools/cyclictest.

system (RTOS), able to provide bounded maximum latencies to high priority real-time user tasks. With the RTOS kernel, the measured latencies while running SPI message exchange never exceeded 457 μs over several hours running at 200 Hz.

2.4.2. Resilient Filesystem

One of the important issues when making reliable Linux embedded systems is how to deal with unexpected power removal. Linux filesystems, in general, are likely to be corrupted if the power is removed while they are performing a write. Even journaling filesystems like ext4 are prone to this. This is why Linux needs to be properly shut down before power is removed, but this is simply not practical for an experimental battery-powered system. Disorderly shutdowns will happen, so this needs to be planned for.

We implement a fully redundant filesystem with error checking using BTRFS (Rodeh et al., 2013) as described in a StackExchange answer.[5] BTRFS is modern journaling filesystem that supports on-the-fly compression and RAID and is capable of self-healing, provided there are redundant copies of the data. The idea is that we create two partitions on the same SD card and mount them as a completely redundant RAID1 array. Any filesystem corruption will be seen as a mismatch between checksum and file, and the redundant copy on the other partition used to replace the corrupt version. This has proven to be very reliable so far, with no corrupted SD cards.

2.4.3. Arena Integration

The Xpucks work within an arena which provides the infrastructure for experiment control, implementing virtual senses if needed, and for logging, see **Figure 6**. It is area 2 m by 1.5 m equipped with a Vicon tracking system and an overhead network webcam. Each Xpuck has a USB WiFi dongle, and the arena has a dedicated WiFi access point. For robustness, each Xpuck has a fixed IP address, and the standard scripts are replaced with a script that continually checks for connectivity to the access point and attempts reconnection if necessary.

Software called the *switchboard* runs on the *Hub* server and is responsible for the distribution of experiments to the Xpucks, their initiation, and the logging of all experiment data. Each Xpuck automatically starts a ROS node at boot which connects to the Hub over ZeroMQ sockets (Hintjens, 2013) supplying a stream of telemetry about the physical state of the Xpuck, including battery levels and power consumption, temperature, sensor values, and actuator settings. The switchboard sends timestamps, virtual sense data, and can command the ROS node to download and execute arbitrary experiment scripts, which would typically set up a more complex ROS graph for the robot controller, which in turn will run the experiment upon a trigger from the switchboard. Controllers are always run locally on the Xpucks. This is all controlled either from the command line on the Hub or with a GUI giving visibility to important telemetry from the swarm.

Each Xpuck is marked with a unique pattern of reflectors recognized by the Vicon system. There are four reflectors arranged on a 4×4 grid with spacing of 10 mm. We used a brute force search to

[5]Corruption-proof SD card filesystem for embedded Linux? http://unix. stackexchange.com/questions/136269/corruption-proof-sd-card-filesystem-for-embedded-linux.
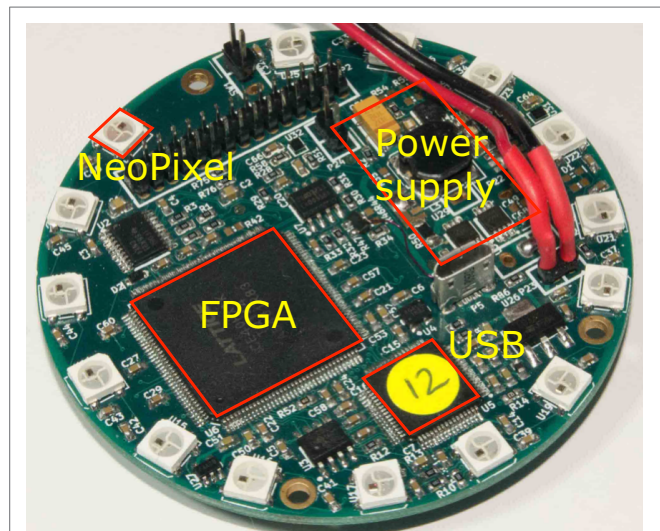
FIGURE 5 | Interface board PCB, showing the boost converter PSU for the XU4 5v supply, the FPGA and USB interface, the VGA camera and SPI level shifting, and the 15 Neopixels.

find unique patterns for each of the 16 Xpucks. Because of the size of the marker pattern and of the Xpucks themselves, there should be no ambiguous conditions when Xpucks are close to each other. This has proved effective with unambiguous detection even when all 16 Xpucks were placed packed together in the arena.

The switchboard software connects to the Vicon system and receives pose messages at the update rate of 50 Hz. This is used to log the absolute positions of the Xpucks during experiments and also to synthesize virtual senses included in the outgoing streams of data from the switchboard to the Xpucks. Range and bearing is an important sense in swarm robotics experiments, which we can construct directly using the e-pucks IR proximity sensors or with additional hardware (Gutiérrez et al., 2009a,b). We can also synthesize range and bearing information from the Vicon data with behavior defined by a *message distribution model*, which allows us to specify parameters such as range, noise, and directionality. There is the capability for Xpucks to send broadcast messages consisting of their ID, this is disseminated by the switchboard according to the message distribution model. Messages received have no content, but are an indication that the sender and the receiver can communicate, actual data transfer can take place point-to-point. In this, we take inspiration from O'Dowd et al. (2014), who use IR communication between e-pucks to establish if contact is possible, data transfer then taking place over WiFi.

2.5. GPGPU Robot Simulator

In this section, we describe the design and realization of a fast parallel physics-based 2D multi robot simulator running on the Xpuck SoC GPU.

To perform onboard evolution of controllers or to evaluate multiple what-if scenarios, we need to be able to run many simulations much faster than real-time. A typical evolutionary algorithm might have a population of p potential solutions. Each of these needs to be evaluated for fitness by running r simulations with different starting conditions. Many generations g of evaluation,

FIGURE 6 | The Xpuck arena. Experiments take place within a 2 m × 1.5 m area surrounded by walls slightly higher than the height of the Xpucks. Each Xpuck has a unique pattern of spherical reflectors on their top surface to enable the Vicon motion tracking system to identify each individuals pose. The Vicon PC is dedicated to managing the Vicon system and makes available a stream of pose data. The Hub PC is responsible for all experiment management, data logging, and virtual sense synthesis.

selection, combination, and mutation take place to produce fitter individuals. Typically, p, r, g might be (50, 10, 100). One scenario we envisage is evolving a controller for the next fixed interval Δt of real time. During the current time interval, we need to complete $n_{sims} = prg$ simulations of that time Δt, or:

$$\frac{n_{sims} \cdot t_{real}}{t_{sim}} < 1 \quad (1)$$

where t_{sim} is the simulated time and t_{real} is the wall clock time for that simulated time. It is generally the case (Vaughan, 2008; Jones et al., 2015) that multi robot simulation time is proportional to the number of agents being simulated. We define a simulator speed using the robot acceleration factor:

$$r_{acc} = \frac{n_{robots} \cdot t_{sim}}{t_{real}} \quad (2)$$

where n_{robots} is the number of robots, t_{sim} and t_{real} as above. With equation (1) we get a required r_{acc} of:

$$r_{acc} > n_{sims} \cdot n_{robots}. \quad (3)$$

We can see that if we are using a single simulator, the required r_{acc} increases with the number of robots being simulated. But if we run a distributed evolutionary algorithm and have a simulator embodied in each robot, the required r_{acc} simply becomes:

$$r_{acc} > n_{sims}. \quad (4)$$

For the example above, we therefore require a simulator with $r_{acc} > 50,000$.

There is a basic trade-off between simulator fidelity and speed. Typical values of r_{acc} when running on a PC are 25 for a full 3D physics simulation like Gazebo, 1,000–2,000 for 2D[6] arbitrary shape simulators with relatively detailed sensory modeling like Stage (Vaughan, 2008), and ARGoS (Pinciroli et al., 2011), and 50,000–100,000[7] for constrained geometry 2D physics game engines like Box2D (Catto, 2009). There is also a cost to generality; the critical path in stage is the ray-tracing operation for modeling of distance sensors, necessary to handle arbitrary object shapes in the simulated world. We show in Jones et al. (2016) that a constrained geometry 2D physics engine simulator is capable of being used to evolve swarm controllers which transfer effectively to the real world, so this motivates our simulator design.

To get good performance on an application running on a GPU, it is necessary that there is a large number of work items that can be performed in parallel. The Mali Midgard GPU architecture present in the Exynos 5422 SoC of the XU4 has six shader cores, each of which can run 256 threads simultaneously. To keep the cores busy, it is recommended that a kernel be executed over hundreds or thousands of work items, depending on its resource usage. We therefore need to design our simulator to have parallelism at least in the hundreds to take advantage of the GPU and be sufficiently constrained in scope that we avoid the costs of generality; by using only straight lines and circles in our simulation, collisions and sensor intersections can be calculated cheaply by geometry, rather than expensive ray-tracing.

---

[6]or "two-and-a-half D" with sensors having some awareness of Z but kinematics and dynamics modelled purely in 2D.
[7]We achieved 80,000 with our Box2D-based kilobot simulator (Jones et al., 2016).

2.5.1. Simulation Model

The simulation models up to 16 Xpuck robots within a 2 m × 1.5 m rectangular arena centered about the origin, with the edges of the arena surrounded with immovable walls. As well as the Xpuck robots, there can be other inert round objects that can be pushed by the Xpucks. The reference model for the robots is given in **Table 3**, this describes the sensors and actuators that are exposed to the robot controller.

We can divide the simulation into three sections; *physics*, *sensing*, and *control*. *Physics* handles the actual physical behavior of the robots within the arena, modeling the dynamics of motion, and collisions in a realistic way. *Sensing* constructs the input variables described in the robot reference model from the locations and attributes of the objects within the simulated world. *Control* runs a robot controller for each simulated robot, responding to the reference model inputs, and producing new output variables, resulting in the robot acting within the world.

There are three types of object within the world: the arena walls, the Xpucks, and inert objects. The walls are immoveable and are positioned horizontally and vertically symmetrically about the origin. Xpucks, which are round and colored, can sense each other with their camera, proximity sensors and range and bearing, and can move with two-wheel kinematics. Inert objects, which are round and colored, can be sensed by Xpuck cameras but not by the proximity sensors because they are low in height. They move only due to collisions.

2.5.1.1. Physics

The physics core of the simulation is based on work by Gaul (2012). There are only circular bodies, which are rigid and have finite mass, and the walls, which have infinite mass. Interactions between bodies are governed by global attributes of coefficients of static and dynamic friction, and restitution. Interactions between the bodies and the arena floor are governed by individual attributes of mass and coefficient of friction. The physical state of each body i is described by the tuple $S_i(\boldsymbol{x}, \boldsymbol{v}, \theta, \omega)$ representing position, velocity, angle, and angular velocity.

The equations of motion governing the system are $\dot{\boldsymbol{v}} = \frac{1}{m}\boldsymbol{F}, \dot{\omega} = \frac{1}{I}\tau, \dot{\boldsymbol{x}} = \boldsymbol{v}, \dot{\theta} = \omega$. They are integrated using the symplectic Euler method (Niiranen, 1999) which has the same computational cost as explicit Euler but better stability and energy preserving properties.

Collisions between bodies are resolved using impulses. For each pair of intersecting bodies, a contact normal and relative velocity are calculated, producing an impulse vector which is used to instantaneously change the linear and angular velocities of the two bodies. This is iteratively applied multiple times to ensure that momentum is transferred in a physically realistic way between multiple contacting bodies.

Collision detection between pairs of bodies with a naive algorithm is $O(n^2)$ so most physics simulators handling a large number of bodies (100 s upwards) use a two stage process with a *broadphase* step that eliminates a high proportion of pairs that cannot possibly be in collision, before the *narrowphase* step that detects and handles those bodies that are actually colliding. But we have only a maximum of 21 bodies (4 walls, 16 robots, 1 object) which means that any broadphase step must be very cheap to actually gain performance overall. We tried several approaches before settling on a simple binning algorithm: each object is binned according to its x coordinate, with bins just larger than the size of the objects. A bin contains a bitmap of all the objects within it. Objects can only be in collision if they are in the same or adjacent bins so the or-combined bitmap of each two adjacent bins is then used to form pairs for detailed collision detection.

The two-wheel kinematics of the robots are modeled by considering the friction forces on each wheel due to its relative velocity to the arena surface caused by the wheel velocity and the object velocity. Friction force is calculated as Coulomb but with μ reduced when the velocity is close to zero using the formulation in Williams et al. (2002): $\mu = \mu_{max}\frac{2 \cdot arctan(k*v)}{\pi}$. With the same justification as Williams et al. (2002), we chose $k = 20$ empirically to ensure numerical stability. The forces on each body are resolved to a single force vector $\boldsymbol{F}$ and torque τ. Non-robot objects simply have zero wheel velocities.

The noise model is a simplified version of that described by Thrun et al. (2005). Three coefficients, $\alpha_1, \alpha_2, \alpha_3$, control, respectively, velocity-dependent position noise, angular velocity-dependent angle noise, and velocity-dependent angle noise. So position and angle are modified: $\boldsymbol{x}' = \boldsymbol{x} + \boldsymbol{v} \cdot s(\alpha_1), \theta' = \theta + \omega \cdot s(\alpha_2) + |\boldsymbol{v}| \cdot s(\alpha_3)$ where $s(\sigma)$ is a sample from a Gaussian distribution with SD σ and mean of zero. Because the noise model is on the critical path of position update and the calculation of even approximate Gaussian noise is expensive, we use a pre-calculated table of random values with the correct distribution.

The physics integration timestep is set at 25 ms for an update rate of 40 Hz. This value was chosen as a trade-off performance and physical accuracy, giving 4 physics steps per controller update cycle.

TABLE 3 | Robot reference model.

Input variables	Values	Description
$P_{i \in \{1,2,\ldots,8\}}$	[0,1]	Proximity sensor i, 0 = nothing in range, 1 = object adjacent to sensor
θ	$[-\pi, \pi)$	Compass, giving pose angle in world frame
$n \in \mathbb{N}$	$\{0, \ldots, 16\}$	Number of neighboring Xpucks
$(r, \angle b)_{i \in \{1, \ldots, n\}, n \neq 0}$	$([r_{min}, r_{max}], [-\pi, \pi))$	Range and bearing of neighbor m
$R_{i \in \{left, center, right\}}$	$\{0, 1\}$	Red blob detection
$G_{i \in \{left, center, right\}}$	$\{0, 1\}$	Green blob detection
$B_{i \in \{left, center, right\}}$	$\{0, 1\}$	Blue blob detection
Output variables		
$v_{i \in \{left, center, right\}}$	$[-v_{max}, v_{max}]$	Left and right wheel velocities
Constants		
t_{update}	100 ms	Controller update period
r_{min}	0.075 m	Minimum range and bearing range (center to center)
r_{max}	0.5 m	Maximum range and bearing range (center to center)
v_{max}	0.13 ms^{-1}	Maximum wheel velocity
$\angle q_{i \in \{1,2,\ldots,8\}}$	0.297, 0.855, 1.571, 2.618, $-2.618, -1.571,$ $-0.855, -0.297$	Angle of proximity sensor i
p_{max}	0.04 m	Proximity sensor maximum range

2.5.1.2. Sensing

There are three types of sensors that need to be modeled. Each Xpuck has eight IR proximity sensors arranged around the body at a height of about 25 mm. These can sense objects out to about 40 mm from the body. The reference model specifies that the reading varies from 0 when nothing is in range, to 1 when there is an object adjacent to the sensor. Similar to the collision detection above, the maximum sensor range is used to set the radius of a circle about the robot which is tested for intersection with other objects. For all cases where there is a possible intersection, a ray is projected from the robot at each sensor angle and a geometrical approximation used to determine the location of intersection with the intersected body and hence the range. This process is actually more computationally expensive than collision detection, but only needs to take place at the controller update rate of 10 Hz.

The second and third types of sensor are the camera blob detection and the range and bearing sense. Blob detection splits the camera field of view into three vertical segments and within each segment, detects the presence of blobs of the primary colors. Range and bearing sense counts the number of robots within 0.5 m and produces a vector pointing to the nearest concentration. Together they are the most computationally expensive of the senses to model. They necessarily traverse the same data structures and so are calculated together.

To model the camera view, we need to describe the field of view subtended by each object, what color it is, and whether it is obscured by nearer objects. We implement this by dividing the visual field into 15 segments and implementing a simple z-buffer. Each object is checked and a left and right extent derived by geometry. The segments that are covered by these extents have the color of the object rendered into them, if the distance to the object is less than that in the corresponding z-buffer entry. As each object is checked, the distance is used to determine if the range and bearing information needs to be updated.

In the real robot arena, range and bearing is implemented as virtual sensing using a Vicon system and communication over WiFi. There is significant latency of around 100–200 ms between a physical position and an updated range and bearing count and vector reaching the real robot controller. Also, the camera on each Xpuck has processing latency of a similar order. For this reason and due to the computational cost, this sensor information is updated at half the controller rate, or 5 Hz.

2.5.1.3. Controller

The controller architecture we use is behavior tree based (Champandard, 2007; Ogren, 2012; Colledanchise and Ogren, 2014; Scheper et al., 2015; Jones et al., 2016). Originating in the games industry for controlling non-player characters, behavior trees are interesting for robotics because they are hierarchical, allowing encapsulation and reuse of sub-behaviors, human readable, aiding analysis of evolved controllers for insight, and amenable to formal analysis. A behavior tree consists of a tree of nodes and a *blackboard* of variables which comprise the interface between the controller and the robot. At every controller update cycle, the tree of each robot is evaluated, with sensory inputs resulting in actuation outputs. Evaluation consists of a depth-first traversal of the tree until certain conditions are met. Each agent has its own blackboard and state memory, the tree is shared by all

agents running the same controller. In our case, we are running homogeneous swarms, so within a particular simulation, only one tree type is used, with each simulated robot running its own instance.

2.5.1.4. Implementation of Simulator on GPU

To best exploit the available performance of the GPU, our implementation must have a high degree of parallelism. We achieve this by running multiple parallel simulations almost entirely within the GPU. The limit to parallelization of running multiple simulations for an evolutionary algorithm is the number of simulations per generation; it is necessary to completely evaluate the fitness of the current generation to create the individuals that will make up the next generation. With the numbers given above, this would be 500 simulations, below what would normally be recommended to keep the GPU busy, but long-lasting threads ensure the GPU is fully utilized.

As we implemented the simulator, it actually turned out that memory organization was the most critical element for performance. Each of the four cores within the first core group of the GPU[8] has a 16-kB L1 data cache and a 256 L2 cache shared between them. Ensuring that data structures for each agent were minimized, and that they fitted within and were aligned to a cache line boundary resulted in large performance improvements. Memory barriers between different stages of the simulation update cycle ensured that data within the caches remained coherent and reduced thrashing. As performance improved and the memory footprint changed, the effect of workgroup size and number of parallel simulations was regularly checked. We used the DS-5 Streamline[9] tool from ARM to visualize the performance counters of the GPU which showed clearly the memory-bound nature of the execution. Profiling of OpenCL applications is difficult at anything finer than the kernel level, so there was much experimentation and whole application benchmarking.

2.6. Image Processing Demonstration

The high computational capability of the Xpuck makes it possible to run camera image processing algorithms not possible on the e-puck on its own or enhanced with the Linux Extension Board. To demonstrate this and to evaluate the performance of the camera, we implement ArUco marker tracking (Garrido-Jurado et al., 2014) and test it with the onboard camera. ArUco is a widely used library that can recognize square black and white fiducial markers in an image and generate camera pose estimations from them. In this demonstration, we use the marker recognition part of the library and test the tracking under different distances and Xpuck rotational velocities.

A ROS node was written to apply the ArUco[10] marker detection library function to the camera image stream and to output the detected ID and pixel coordinates on a ROS topic. Default detection options were used and no particular attention was paid to optimization.

---

[8]The six cores are divided into two core groups, one with four cores and one with two. These are presented as two separate OpenCL devices. For ease of coding, only one core group was used.

[9]https://developer.arm.com/products/software-development-tools/ds-5-development-studio/streamline.

[10]Version 1.2, standard install from Ubuntu 14.04.4 ROS Indigo repository.

TABLE 4 | Speed of simulator with various functionalities enabled.

	256 simulations			512 simulations		
Functionality	t_{rss}	Δt_{rss}	%	t_{rss}	Δt_{rss}	%
Physics	6.9	6.9	40	6.7	6.7	31
+Sensors	11	4.5	26	11	4.0	19
+Camera and R&B	15	3.1	18	14	3.3	16
+Controller (all functionality)	17	2.6	16	21	7.2	34

16 robots, 1 passive object, basic exploration and collision avoidance controller. Tested over five runs with 256 and 512 parallel simulations. t_{rss} is time (μs) per robot simulated second. With 256 parallel simulations, the physics functionality dominates at 40% of the processing time, but with 512 parallel simulations, controller processing is the largest proportion.

Two experiments were conducted. In both cases, we used video from the Xpuck camera at a resolution of 320×240 and a frame rate of 15 Hz. First, we measured the time taken to process an image with the detection function under conditions of no markers, four 100 mm markers in a 2×2 grid, and 81 20 mm markers in a 9×9 grid. Frame times were captured for 60 s.

Second, we affixed four ArUco tags of size 100 mm with different IDs to locations along the arena walls. An Xpuck was placed in three different locations within the arena and commanded to rotate at various speeds up to 0.7 rad/s. Data were collected for 31,500 frames. Commanded rotational velocity, Vicon tracking data, and marker tracking data were all captured for analysis.

The data were analyzed in the following way: each video frame is an observation, which may have markers present within it. Using a simple geometrical model, we predict from the Vicon data and the known marker positions whether a marker should be visible in a given frame and check this against the output of the detector for that frame. From this, we derive detection probability curves for different rotation speeds.

2.7. In-Swarm Evolution Demonstration

One of our motivations for moving computation into the swarm is to tackle the scalability of swarm controller evolution. To demonstrate both the computational capability of the Xpuck swarm and scalability, we implement an island model evolutionary algorithm and demonstrate performance improvement when running on multiple Xpuck robots.

The island model of evolutionary algorithms divides the population of individuals into multiple subpopulations, each of which follows its own evolutionary trajectory, with the addition of *migration*, where some individuals of the subpopulations are shared or exchanged with other subpopulations. Island model evolutionary algorithms enable coarse-grained parallelism, with each island corresponding to a different compute node, and sometimes outperform single population algorithms by maintaining diversity (Whitley et al., 1999). Even without that factor, the ability to scale the size of total population with the number of compute nodes hosting subpopulations is desirable for a swarm of robots running embodied evolution.

2.7.1. Implementation of Island Model

On each Xpuck, we run a genetic algorithm evolving a population of behavior tree controllers similar to that described in Jones et al. (2016) using methods from Genetic Programming (Koza, 1992). The parameters are described in **Table 5**. Evolution proceeds as follows: an initial subpopulation of n_{sub} individuals is generated using the Koza's *ramped_half_and_half* procedure, detailed in Poli et al. (2008), with a depth of n_{depth}. Each individual is evaluated for fitness by running n_{sims} simulations with different starting conditions and averaging the individual fitnesses. The subpopulation is sorted and the top n_{elite} individuals are copied unchanged into the new subpopulation. The remaining slots are filled by tournament selection of two individuals with replacement followed by a tree crossover operation, with random node selection biased to internal nodes 90% of the time (Koza, 1992), to create a new individual. Then, every parameter within that individual is mutated with probability p_{mparam}, followed by mutating every node to another of the same arity with probability p_{mpoint}, followed by replacing a subtree with a new random subtree with probability $p_{msubtree}$. This new population is then used for the next round of fitness measurement.

The genetic algorithm is extended to the island model in the following way: after every n_{epoch} generations, each Xpuck sends a copy of the fittest individual in its subpopulation to its neighbors. They replace the weakest individuals in their subpopulations. Currently, this is mediated through a *genepool server*, running on the Hub PC, although direct exchange of genetic material between individual Xpucks is also possible using local IR communication. This server maintains the topology and policy for connecting the islands. This may be physically based, drawing on the position information from the Vicon. It is important to note that server provides a way to abstract and virtualize the migration of individuals; in the same way, we use the Vicon information to provide virtual sensing. When the server receives an individual from a node, it replies with a set of individuals, according to the policy. These are used to replace the least fit individuals on the requesting node. The process is asynchronous, not requiring that the nodes execute generations in lockstep. The policy for this experiment is to make a copy of each individual available to every other node, so with n_{nodes} nodes the migration rate is $r_{migration} = \frac{n_{nodes}-1}{n_{sub} \cdot n_{epoch}}$.

2.7.2. Task and Fitness Function

We evolve a behavior tree controller for a collective object movement task. The task takes place in a $2\,\text{m} \times 1.5\,\text{m}$ arena with the origin at the center and surrounded by walls greater than the height of the Xpucks. The walls and floor are white. A blue plastic frisbee of 220 mm diameter is placed at the origin. Nine Xpucks with red skirts are placed in a grid with spacing 100 mm centered at $(-0.8, 0)$ and facing rightwards. The goal is to push the frisbee to the left. Fitness is based on how far to the left the frisbee is after a fixed time. An individual Xpuck can push the frisbee, but only at about half the full Xpuck speed, so collective solutions

have the potential to be faster. The swarm is allowed to execute its controller for 30 s. After this time, the fitness is given by equation (5).

$$f = \begin{cases} r_{derate} \frac{-x}{1-l_{frisbee\_radius}}, & \text{for } x < 0 \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

where x is the x-coordinate of the center of the frisbee, and r_{derate} is a means of bloat control, proportionately reducing the fitness of behavior trees which use more than 50% of the resources available. To show scalability with increasing numbers of Xpucks, we compare two scenarios, first a single Xpuck running a standalone evolution and second six Xpucks running an island model evolution. In both cases, the parameters are as in **Table 5**. With the island model, every $n_{epoch} = 2$ generations, a node sends to all its neighbors a copy of its fittest individual and receives their fittest individuals, using these to replace its five least fit individuals, giving a migration rate $r_{migration} = 0.078$. Each scenario is run ten times with different initial random seeds.

3. RESULTS

3.1. Xpucks

The total cost of 25 Xpucks was £3,325, or £133 each. This includes all parts, PCBs, XU4 single board computers, and batteries. It does not include assembly or the base e-pucks, which cost around £700. Although it should be possible for a university technician to assemble the boards in small quantities, the approximate costs per board for factory PCB assembly were £17 for 25 boards, dropping rapidly to £6 for 100 boards.[11] It is our intention to make the design open source and freely available.[12]

Currently, we have 16 assembled and functional robots. Battery life when running a moderate computational load is close to 3 h. When continuously running the extremely computationally demanding evolutionary algorithm described, the battery life dropped to around 1 h 20 min.

3.2. Simulator

Table 4 shows the results of running parallel simulations for a simulated time of 30 s. Each simulation consists of 16 robots running a simple controller for exploration with basic collision avoidance, and one additional object that can be pushed by the

---

[11]Online quote from https://www.pcbway.com/.
[12]https://bitbucket.org/siteks/xpuck_design.

TABLE 5 | Evolutionary algorithm parameters.

Parameter	Value	Description
n_{gens}	100	Number of generations
n_{sub}	32	Size of subpopulation
n_{sims}	8	Number of simulations for evaluation of fitness
n_{elite}	3	Size of elite
p_{mparam}	0.05	Probability of mutating a parameter
p_{mnode}	0.05	Probability of mutating a node
$p_{msubtree}$	0.05	Probability of replacing a subtree
n_{depth}	4	Maximum depth of tree generated
n_{tsize}	3	Tournament size
n_{epoch}	2	Migration epoch

robots. The effect of running different numbers of parallel scenes and with various different levels of functionality enabled is shown. t_{rss} is the time to simulate one robot second. $t_{rss} = \frac{1}{r_{acc}}$, so the required acceleration factor of 50,000 corresponds to $t_{rss} = 20\ \mu s$. It can be seen that the requirement is met when running 256 simulations in parallel, with $t_{rss} = 17\ \mu s$. It is interesting to note that when running 512 simulations, the performance is better with all functionalities except the controller enabled. We surmise that, when running the controller, the total working set is such that there is increased cache thrashing with 512 parallel simulations.

The performance of the simulator running on the Xpuck GPU is comparable to the same code running on the CPU of a much more powerful desktop system and at least ten times faster than more general purpose robot simulators such as Stage and ARGoS running on the desktop. Although future work will aim to demonstrate the transferability of the evolved solutions, we note that the fidelity of the simulator is similar to previous work (Jones et al., 2016) which successfully transferred with only moderate reality gap effects.

3.3. Image Processing

For the computationally demanding image processing task, **Table 6** shows the time taken for the Xpuck to process a 320 × 240 pixel frame using the ArUco library to search for markers. With four large markers, the 23 ms processing time is fast enough to sustain the full camera frame rate of 15 Hz. In the 81 marker case, detection speed slows to 94 ms, such that a 15 Hz rate is not sustainable. In both cases, however, all the markers were correctly detected in each frame.

The dsPIC of the e-puck would not be capable of running this code—it is only capable of capturing camera video at 40 × 40 pixels and 4 Hz with no image processing (Mondada et al., 2009) and has insufficient RAM to hold a complete image. The Linux Extension Board processor could potentially run the detection code, but we estimate the processing time would be at least 50 times longer[13] giving a frame rate of less than 1 Hz.

The arena detection experiment collected 31,500 frames, with 11,076 marker detections possible in ideal circumstances. Actual detections numbered 8,947, a total detection rate of 81%. **Figure 7** shows the probability of detecting a marker under different conditions. With four markers around the arena, and the Xpuck capturing data at three locations within the arena, there are twelve distance/angle combinations. Distances vary from 0.5 to 1.5 m, and angles from 0° to 70°. The gray envelope and lines show the individual distance/angle combinations against the angular

---

[13]ARM926EJS @200 MHz = 220DMIPS, A15 @800 MHz = 2800DMIPS, 4× penalty for no floating point, single core only: 50×.

TABLE 6 | ArUco detector speed at a resolution of 320 × 240 pixels under different conditions.

Condition	Processing time (ms)	σ
No markers	12.4	2.7
4 mm × 100 mm markers	23.3	4.5
81 mm × 20 mm markers	93.8	0.25

In each case, the input was for 60 s. The detector code is unable to process frames at the full 15 Hz in the 81 marker case.
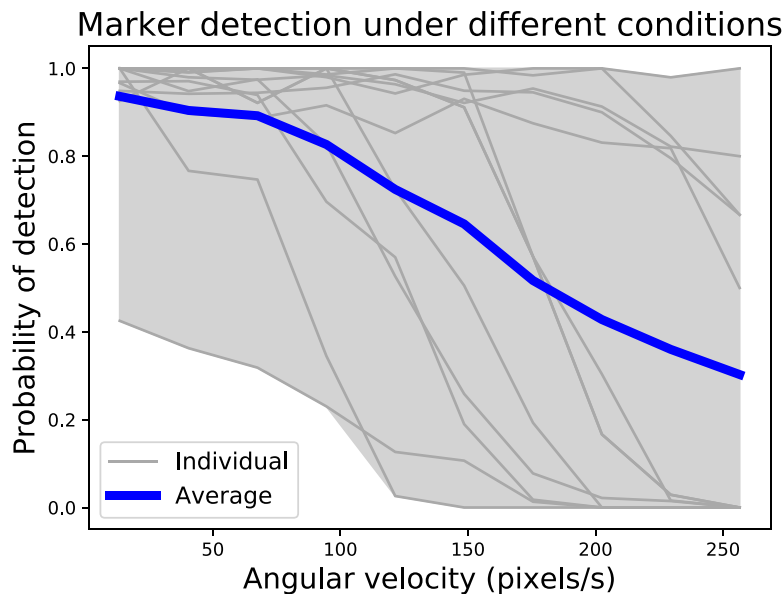
FIGURE 7 | Probability of marker detection under different conditions. There are four markers around the arena, with data collected at three locations, giving twelve distance/angle combinations. Observations at a resolution of 320 × 240 pixels were made for 31,500 frames, with 8,947 marker detections out of a possible 11,076, a detection rate of 81%. The number of detections compared to the maximum possible for each geometry was binned by angular velocity to give probability curves. Gray lines are individual distance/angle combinations, and the blue line is the average over all combinations. Generally, detection rate falls with increasing angular velocity, with a 50% detection rate at 180 pixels/s.

velocity, with the blue line being the average over all observations. Angular velocity is expressed in pixels/s for better intuition about how fast a marker is traversing the field of view of the camera. Generally, the detection rate falls as the angular velocity increases, with a 50% detection rate at 180 pixels/s.

This shows that, even with unoptimized code, the Xpuck has sufficient computational performance, and the camera subsystem is of sufficient quality, that visual marker tracking is feasible.

3.4. Evolution

The results are summarized in **Figure 8**. It is clear that the six node island model evolutionary system performs better than the single node. Maximum fitness reached is higher at 0.7 vs 0.5, and progress is faster. Of interest is the very low median fitness of the single node populations (shown with red bar in boxes), compared to the mean. This is because seven out of the ten runs never reached a higher fitness than 0.1 suggesting the population size or the number of generations is too small. Conversely, the median and mean of the island model population's maximum fitnesses are quite similar, showing a more consistent performance across runs. If we look at how fast the mean fitness rises, a single node takes 100 generations for the fitness to reach 0.15. The six node system reaches this level of mean fitness after 25 generation, four times faster.

Figure 9 shows a plot of the elapsed processing time per generation over ten runs. The variation is mostly due to the complexity and depth of the behavior tree controllers within each generation, together with the trajectory of the robots in simulation. Each of the ten runs of both the island model and the single node systems completed in less than 10 min. For comparison, each evolutionary run in our previous work (Jones et al., 2016) took several hours on a powerful desktop machine.

This demonstrates the Xpucks are sufficiently capable to host in-swarm evolutionary algorithms that scale in performance with the size of the swarm.

4. DISCUSSION

4.1. Background and Related Work

In the introduction, we outline three areas which we feel could benefit from the increased processing power of the Xpuck.

Swarm robotics (Sahin, 2005) takes inspiration from collective phenomena in nature, where global behaviors emerge from the local interactions of the agents of the swarm with each other, and with the environment. The design of controllers such that a desired collective behavior emerges is a central problem. Common approaches use bioinspiration, evolution, reverse engineering, and hand-design (Reynolds, 1987; Trianni et al., 2003; Hauert et al., 2009b; Trianni and Nolfi, 2011; Francesca et al., 2014). The controller architectures include neural networks, probabilistic finite state machines, behavior trees, and hybrid combinations (Baldassarre et al., 2003; Francesca et al., 2015; Duarte et al., 2016; Jones et al., 2016). See Francesca and Birattari (2016) for a recent review. When using evolution or other methods of automatic design within an off-line simulated environment, the problem of the transferability of the controller from simulation to real robots arises, the so-called *reality gap*. There are various approaches to alleviating this such as noise injection within a minimal simulation (Jakobi et al., 1995; Jakobi, 1998), making transferability a goal within the evolutionary algorithm (Koos et al., 2013; Mouret and Chatzilygeroudis, 2017), and reducing the representational power of the controller (Francesca et al., 2014, 2015). Embodied evolution directly tests candidate controllers in reality. When applied to swarms (Watson et al., 2002) the evolutionary algorithm
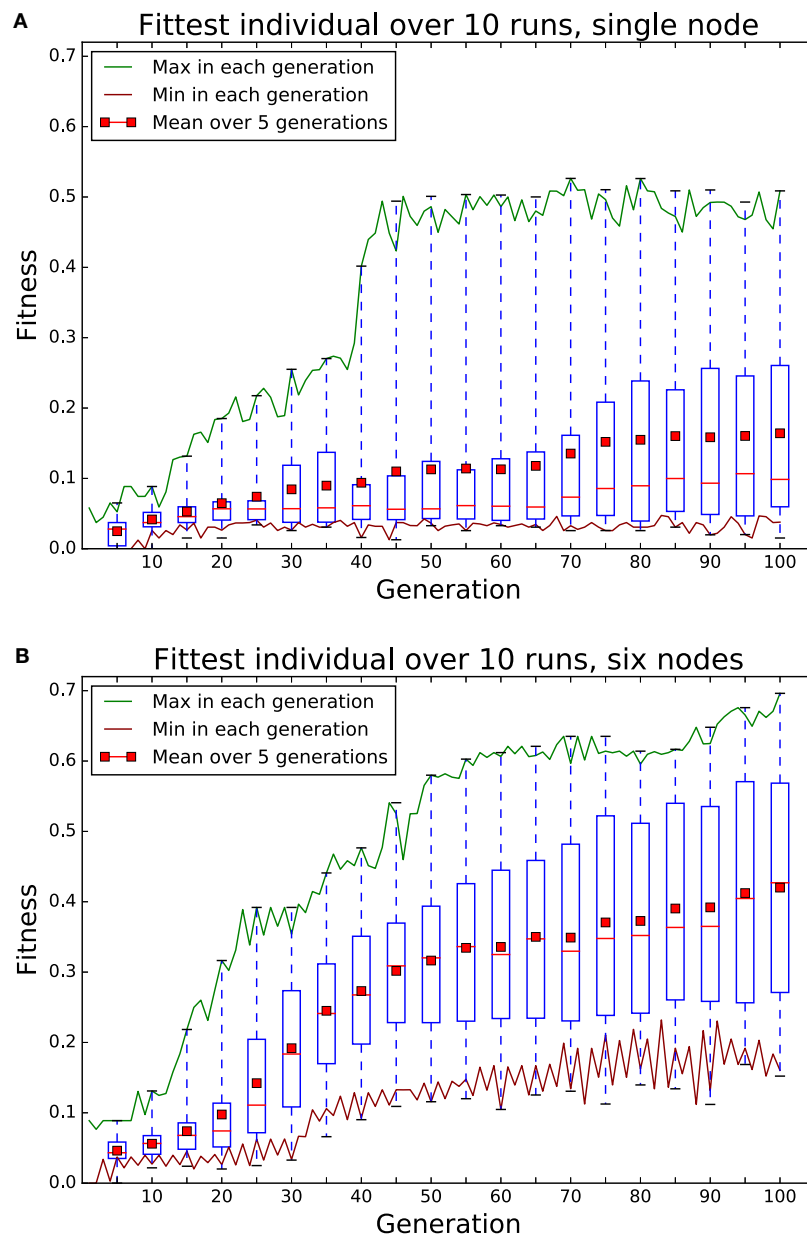
FIGURE 8 | Comparison of 100 generations of evolution using a single node **(A)** and using an island model with six nodes **(B)**. Each node has a population of 32 individuals, evaluated 8 times with different starting conditions for fitness. Each node in the six node system replaces is five least fit individuals with the fittest from the other five nodes every two generations. Boxes summarize data for that generation and the previous four. Red bar in boxes indicates median. The six node system clearly shows higher maximum fitness after 100 generations and reaches the same mean fitness as the single node system in a quarter of the time. The large difference between mean and median in the single node system is due to seven of the ten runs not exceeding a fitness of 0.1.

is distributed over the robots (Takaya and Arita, 2003; Bredeche et al., 2012; Doncieux et al., 2015). Other approaches use reality sampling to alter the simulated environment to better match true fitnesses (Zagal et al., 2004; O'Dowd et al., 2014). This requires either off-board processing with communication links to the robot or sufficient processing power on the robot to run simulations. Related is the concept of surrogate fitness functions (Jin, 2011) with cheap but inaccurate fitness measures made in simulation and expensive but accurate measures made in reality.

Using internal simulation, models can be means of detecting malfunction and adapting (Bongard et al., 2006), or asking *what-if*

questions, so as to evaluate the consequences of possible actions in simulation (Marques and Holland, 2009). This is applied to the fields of both robot safety and machine ethics in Winfield et al. (2014), Winfield (2015), Blum et al. (2018), and Vanderelst and Winfield (2018). It is obvious that any robot relying on simulation for its ethical or safe behavior must embody that simulation and not use unreliable communications links. Swarms are usually assumed to be robust to failure, but Bjerknes and Winfield (2013) show that this is not always the case. By using internal models and observing other agents within the swarm, agents not behaving as predicted can be identified (Millard et al., 2013, 2014).

FIGURE 9 | Time per generation of single node evolution, 10 runs of 100 generations each, with different starting conditions. The average length of a run is 5.7 min. The variation in processing time is due mostly to the size and complexity of the behavior trees within the population.

The social insects that are often the inspiration for swarm robotics are actually far more complex than the commonly used ANN controllers of swarm robot agents. They have many more neurons, and the neurons are behaviorally complex. The computational requirement of simulating biologically plausib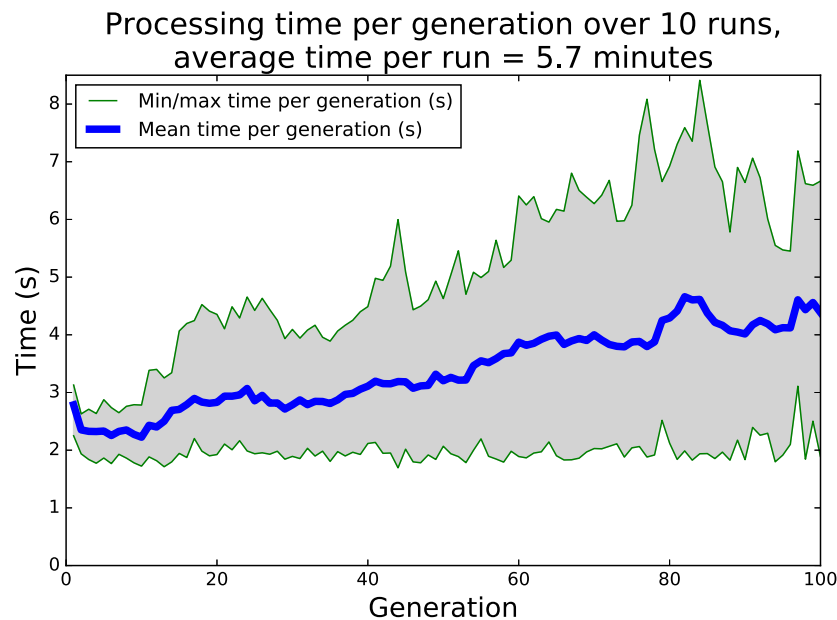le neurons can be estimated. The Izhikevich (2003) model is commonly used and reported performances vary between 7 and 50 MFLOPS/neuron (Ananthanarayanan et al., 2009; Fidjeland and Shanahan, 2010; Scorcioni, 2010; Minkovich et al., 2014). The system we describe could plausibly simulate several thousand biologically plausible neurons per Xpuck.

A number of different platforms have been used for swarm robotics research. The e-puck by Mondada et al. (2009) is widely used for experiments with numbers in the tens. Rubenstein et al. (2012) introduced the Kilobot, which enables swarm experiments involving over 1,000 low-cost robots. Both platforms work on a 2D surface. Other platforms include Swarmbots (Dorigo et al., 2004), R-one (McLurkin et al., 2013), and Pheeno (Wilson et al., 2016). Swarm platforms working in 3D are also described, Hauert et al. (2009a) demonstrate Reynolds flocking (Reynolds, 1987) with small fixed-wing drones, see also Kushleyev et al. (2013) and Vásárhelyi et al. (2014). Most described platforms are homogeneous, but heterogeneous examples exist such as the Swarmanoid (Dorigo et al., 2013). **Table 1** compares some of these platforms, looking at cost and processing power. It is only with the very recent platforms of the Pi-puck and Pheeno (unavailable at the time of design of Xpuck) that the processing power exceeds 1.2 GFLOPS.

We designed the Xpuck explicitly with the e-puck in mind, because, like many labs, we already have a reasonably large number of them. The e-puck is very successful, with in excess of 3,500 shipped units, perhaps due to its simple reliable design and

extendability. Expansion connectors allow additional boards that add capabilities. Three such are relevant here because they extend the processing power of the e-puck. The Linux Extension Board (Liu and Winfield, 2011) adds a 200-MHz Atmel ARM processor running embedded Linux, with WiFi communication. The e-puck extension for Gumstix Overo COM is a board from GCTronic that interfaces a small Linux single board computer, the Gumstic Overo Earthstorm,[14] to the e-puck. A recent addition is the Pi-puck (Millard et al., 2017) which provides a means of using the popular Raspberry Pi single board computers to control an e-puck. The extension board connects the Pi to the various interfaces of the e-puck and provides full access to all sensors and actuators except the camera.

4.2. Conclusion

We have presented the Xpuck swarm, a new research platform with an aggregate raw processing power in excess of two Teraflops. The swarm of 16 e-puck robots augmented with custom hardware uses the substantial CPU and GPU processing power available from modern mobile System-on-Chip devices; each individual Xpuck has at least an order of magnitude greater compute performance than previous swarm robotics platforms. As well as the robots themselves, we have described the system as a whole that allows us to run new classes of experiments that require high-individual robot computation and large numbers of robots. We foresee many uses such as online evolution or learning of swarm controllers, simulation of what-if questions about possible actions, distributed super-computing for mobile platforms, and real-world applications of swarm robotics that requires image processing, or distributed SLAM.

---

[14]https://store.gumstix.com/coms/overo-coms/overo-earthstorm-com.html.

To demonstrate the capabilities of the system, we have shown the feasibility of running a widely used fiducial marker recognition image processing library, which could form the basis for a distributed swarm localization system. We have implemented a fast robot simulator tailored specifically to run on the GPU of the Xpuck. The performance of this simulator on the Xpuck GPU is comparable to the same code running on the CPU of a much more powerful desktop system, and at least ten times faster than general purpose simulators such as Stage and ARGoS running on the desktop. By using this fast simulator within an island model evolutionary algorithm, we have demonstrated the ability to perform in-swarm evolution. The increasing performance at reaching a given fitness with increasing Xpuck swarm size demonstrates the scalability of this approach. Previous work of ours used evolutionary algorithms that took several hours on the desktop to achieve what is now possible in less than 10 min on the swarm.

In conclusion, we present a new tool for investigating collective behaviors. Our platform provides vastly increased computational performance situated within the swarm itself, opening up the possibility of novel approaches and algorithms.

AUTHOR CONTRIBUTIONS

SJ substantial contributions to the conception of the work, design and implementation of robots, design and implementation of software, design of experiments, acquisition and analysis of data, drafting the work, final approval, and agreement to be accountable. MS, SH, and AW substantial contributions to the conception of the work, critical revision of the work, final approval, and agreement to be accountable.

FUNDING

REFERENCES

Ananthanarayanan, R., Esser, S. K., Simon, H. D., and Modha, D. S. (2009). "The cat is out of the bag: cortical simulations with 10^9 neurons, 10^{13} synapses," in *High Performance Computing Networking, Storage and Analysis, Proceedings of the Conference on* (New York, NY: IEEE), 1–12.

Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolving mobile robots able to display collective behaviors. *Artif. Life* 9, 255–267. doi:10.1162/106454603322392460

Bjerknes, J. D., and Winfield, A. F. (2013). "On fault tolerance and scalability of swarm robotic systems," in *Distributed Autonomous Robotic Systems* (Berlin: Springer), 431–444.

Blum, C., Winfield, A. F. T., and Hafner, V. V. (2018). Simulation-based internal models for safer robots. *Front. Robot. AI* 4:74. doi:10.3389/frobt.2017.00074

Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science* 314, 1118–1121. doi:10.1126/science.1133687

Bredeche, N., Montanier, J.-M., Liu, W., and Winfield, A. F. (2012). Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. *Math. Comput. Model. Dyn. Syst.* 18, 101–129. doi:10.1080/13873954.2011.601425

Catto, E. (2009). *Box2D: A 2D Physics Engine for Games*. Available from: http://box2d.org/about/

Champandard, A. (2007). "Behavior trees for next-gen game AI," in *Game Developers Conference, Audio Lecture*, Lyon.

Colledanchise, M., and Ogren, P. (2014). "How behavior trees modularize robustness and safety in hybrid systems," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on* (Chicago, IL: IEEE), 1482–1488.

Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. E. G. (2015). Evolutionary robotics: what, why, and where to. *Front. Robot. AI* 2:4. doi:10.3389/frobt.2015.00004

Dorigo, M., Floreano, D., Gambardella, L. M., Mondada, F., Nolfi, S., Baaboura, T., et al. (2013). Swarmanoid: a novel concept for the study of heterogeneous robotic swarms. *IEEE Robot. Autom. Mag.* 20, 60–71. doi:10.1109/MRA.2013.2252996

Dorigo, M., Tuci, E., Groß, R., Trianni, V., Labella, T. H., Nouyan, S., et al. (2004). "The swarm-bots project," in *International Workshop on Swarm Robotics* (Berlin: Springer), 31–44.

Duarte, M., Gomes, J., Costa, V., Oliveira, S. M., and Christensen, A. L. (2016). *Hybrid Control for a Real Swarm Robotics System in an Intruder Detection Task*. Berlin: Springer International Publishing, 213–230.

Fidjeland, A. K., and Shanahan, M. P. (2010). "Accelerated simulation of spiking neural networks using GPUs," in *Neural Networks (IJCNN), The 2010 International Joint Conference on* (Barcelona: IEEE), 1–8.

Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Curr. Biol.* 17, 514–519. doi:10.1016/j.cub.2007.01.058

Francesca, G., and Birattari, M. (2016). Automatic design of robot swarms: achievements and challenges. *Front. Robot. AI* 3:29. doi:10.3389/frobt.2016.00029

Francesca, G., Brambilla, M., Brutschy, A., Garattoni, L., Miletitch, R., Podevijn, G., et al. (2015). AutoMoDe-chocolate: automatic design of control software for robot swarms. *Swarm Intell.* 9, 125–152. doi:10.1007/s11721-015-0107-9

Francesca, G., Brambilla, M., Brutschy, A., Trianni, V., and Birattari, M. (2014). AutoMoDe: a novel approach to the automatic design of control software for robot swarms. *Swarm Intell.* 8, 89–112. doi:10.1007/s11721-014-0092-4

Garrido-Jurado, S., Noz Salinas, R. M., Madrid-Cuevas, F., and Marín-Jiménez, M. (2014). Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognit.* 47, 2280–2292. doi:10.1016/j.patcog.2014.01.005

Gaul, R. (2012). *Impulse Engine 2D Physics Simulator*. Available from: http://www.randygaul.net/projects-open-sources/impulse-engine/

Gronqvist, J., and Lokhmotov, A. (2014). "Optimising OpenCL kernels for the ARM Mali-T600 GPUs," in *GPU Pro 5: Advanced Rendering Techniques* (Boca Raton, FL: CRC Press), 327–357.

Gurrum, S. P., Edwards, D. R., Marchand-Golder, T., Akiyama, J., Yokoya, S., Drouard, J.-F., et al. (2012). "Generic thermal analysis for phone and tablet systems," in *Electronic Components and Technology Conference (ECTC), 2012 IEEE 62nd* (New York, NY: IEEE), 1488–1492.

Gutiérrez, Á, Campo, A., Dorigo, M., Donate, J., Monasterio-Huelin, F., and Magdalena, L. (2009a). "Open e-puck range & bearing miniaturized board for local communication in swarm robotics," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (Kobe: IEEE), 3111–3116.

Gutiérrez, Á, Tuci, E., and Campo, A. (2009b). Evolution of neuro-controllers for robots' alignment using local communication. *Int. J. Adv. Robot. Syst.* 6, 6. doi:10.5772/6766

Hauert, S., Zufferey, J.-C., and Floreano, D. (2009a). Evolved swarming without positioning information: an application in aerial communication relay. *Auton. Robots* 26, 21–32. doi:10.1007/s10514-008-9104-9

Hauert, S., Zufferey, J.-C., and Floreano, D. (2009b). "Reverse-engineering of artificially evolved controllers for swarms of robots," in *Evolutionary Computation, 2009. CEC'09. IEEE Congress on* (New York, NY: IEEE), 55–61.

Hintjens, P. (2013). *ZeroMQ: Messaging for Many Applications*. Sebastopol, CA: O'Reilly Media, Inc.

Ho, J., and Smith, R. (2015). *NVIDIA Tegra X1 Preview and Architecture Analysis*. Available at: http://www.anandtech.com/show/8811/nvidia-tegra-x1-preview

Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Trans. Neural Networks* 14, 1569–1572. doi:10.1109/TNN.2003.820440

Jakobi, N. (1998). "Running across the reality gap: Octopod locomotion evolved in a minimal simulation," in *Evolutionary Robotics* (Berlin: Springer), 39–58.

Jakobi, N., Husbands, P., and Harvey, I. (1995). "Noise and the reality gap: the use of simulation in evolutionary robotics," in *Advances in Artificial Life* (Berlin: Springer), 704–720.

Jin, Y. (2011). Surrogate-assisted evolutionary computation: recent advances and future challenges. *Swarm Evol. Comput.* 1, 61–70. doi:10.1016/j.swevo.2011.05.001

Jones, S., Studley, M., Hauert, S., and Winfield, A. (2016). "Evolving behaviour trees for swarm robotics," in *Springer Tracts in Advanced Robotics: 13th International Symposium on Distributed Autonomous Robotic Systems (DARS 2016)*, London, UK.

Jones, S., Studley, M., and Winfield, A. (2015). "Mobile GPGPU acceleration of embodied robot simulation," in *Artificial Life and Intelligent Agents: First International Symposium, ALIA 2014* (Cham: Bangor, UK). Revised Selected Papers, Communications in Computer and Information Science. Springer International Publishing.

Khronos OpenCL Working Group. (2010). *The OpenCL Specification, Version 1.1.* Beaverton, OR: Khronos Group. Available at: https://www.khronos.org/registry/OpenCL/specs/opencl-1.1.pdf

Koos, S., Mouret, J.-B., and Doncieux, S. (2013). The transferability approach: crossing the reality gap in evolutionary robotics. *Evol. Comput. IEEE Trans.* 17, 122–145. doi:10.1109/TEVC.2012.2185849

Koza, J. R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, Vol. 1. Cambridge, MA: MIT Press.

Kushleyev, A., Mellinger, D., Powers, C., and Kumar, V. (2013). Towards a swarm of agile micro quadrotors. *Auton. Robot.* 35, 287–300. doi:10.1007/s10514-013-9349-9

Liu, W., and Winfield, A. F. (2011). Open-hardware e-puck Linux extension board for experimental swarm robotics research. *Microprocess. Microsyst.* 35, 60–67. doi:10.1016/j.micpro.2010.08.002

Mack, C. A. (2011). Fifty years of Moore's law. *IEEE Trans. Semicond. Manuf.* 24, 202–207. doi:10.1109/TSM.2010.2096437

Marques, H. G., and Holland, O. (2009). Architectures for functional imagination. *Neurocomputing* 72, 743–759. doi:10.1016/j.neucom.2008.06.016

McLurkin, J., Lynch, A. J., Rixner, S., Barr, T. W., Chou, A., Foster, K., et al. (2013). "A low-cost multi-robot system for research, teaching, and outreach," in *Distributed Autonomous Robotic Systems* (Berlin: Springer), 597–609.

Menzel, R., and Giurfa, M. (2001). Cognitive architecture of a mini-brain: the honeybee. *Trends Cogn. Sci.* 5, 62–71. doi:10.1016/S1364-6613(00)01601-6

Millard, A. G., Joyce, R. A., Hilder, J. A., Fleseriu, C., Newbrook, L., Li, W., et al. (2017). "The Pi-puck extension board: a Raspberry Pi interface for the e-puck robot platform," in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (York, UK: York).

Millard, A. G., Timmis, J., and Winfield, A. F. (2013). "Towards exogenous fault detection in swarm robotic systems," in *Conference towards Autonomous Robotic Systems* (Berlin: Springer), 429–430.

Millard, A. G., Timmis, J., and Winfield, A. F. (2014). "Run-time detection of faults in autonomous mobile robots based on the comparison of simulated and real robot behaviour," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on* (Chicago, IL: IEEE), 3720–3725.

Minkovich, K., Thibeault, C. M., O'Brien, M. J., Nogin, A., Cho, Y., and Srinivasa, N. (2014). HRLSim: a high performance spiking neural network simulator for GPGPU clusters. *IEEE Trans. Neural Networks Learn. Syst.* 25, 316–331. doi:10.1109/TNNLS.2013.2276056

Mitri, S., Floreano, D., and Keller, L. (2009). The evolution of information suppression in communicating robots with conflicting interests. *Proc. Natl. Acad. Sci. U.S.A.* 106, 15786–15790. doi:10.1073/pnas.0903152106

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., et al. (2009). "The e-puck, a robot designed for education in engineering," in *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*, Castelo Branco, Vol. 1, 59–65.

Mouret, J.-B., and Chatzilygeroudis, K. (2017). "20 years of reality gap: a few thoughts about simulators in evolutionary robotics," in *Workshop "Simulation in Evolutionary Robotics", Genetic and Evolutionary Computation Conference*, New York, NY.

Niiranen, J. (1999). "Fast and accurate symmetric Euler algorithm for electromechanical simulations NOTE: the method became later known as "Symplectic Euler"," in *Proceedings of the 6th International Conference ELECTRIMACS '99:*

Modelling and Simulation of Electric Machines, Converters and Systems, Vol. 1 (Lisbon: Lisboa, Portugal), 71–78.

Nvidia. (2007). *NVIDIA CUDA, Compute Unified Device Architecture Programming Guide*. Santa Clara, CA: NVIDIA.

O'Dowd, P. J., Studley, M., and Winfield, A. F. (2014). The distributed co-evolution of an on-board simulator and controller for swarm robot behaviours. *Evol. Intell.* 7, 95–106. doi:10.1007/s12065-014-0112-8

Ogren, P. (2012). "Increasing modularity of UAV control systems using computer game behavior trees," in *AIAA Guidance, Navigation and Control Conference* (Minneapolis, MN).

Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., et al. (2011). "ARGoS: a modular, multi-engine simulator for heterogeneous swarm robotics," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (San Francisco, CA: IEEE), 5027–5034.

Poli, R., Langdon, W. B., McPhee, N. F., and Koza, J. R. (2008). *A Field Guide to Genetic Programming*. Available at: http://www.gp-field-guide.org.uk

Polilov, A. A. (2012). The smallest insects evolve anucleate neurons. *Arthropod Struct. Dev.* 41, 29–34. doi:10.1016/j.asd.2011.09.001

Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., et al. (2009). "ROS: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, Vol. 3 (Kobe, Japan), 5.

Reynolds, C. W. (1987). "Flocks, herds and schools: a distributed behavioral model," in *ACM SIGGRAPH Computer Graphics* (New York, NY: ACM), 21, 25–34.

Rodeh, O., Bacik, J., and Mason, C. (2013). BTRFS: the Linux B-tree filesystem. *ACM Trans. Storage* 9, 9. doi:10.1145/2501620.2501623

Rostedt, S., and Hart, D. V. (2007). "Internals of the RT patch," in *Proceedings of the Linux Symposium*, Ottawa, Vol. 2, 161–172.

Rubenstein, M., Ahler, C., and Nagpal, R. (2012). "Kilobot: a low cost scalable robot system for collective behaviors," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (St Paul, MA: IEEE), 3293–3298.

Sahin, E. (2005). "Swarm robotics: from sources of inspiration to domains of application," in *Swarm Robotics* (Berlin: Springer), 10–20.

Scheper, K. Y., Tijmons, S., de Visser, C. C., and de Croon, G. C. (2016). Behavior trees for evolutionary robotics. *Artif. Life.* 22, 23–48. doi:10.1162/ARTL_a_00192

Scorcioni, R. (2010). "GPGPU implementation of a synaptically optimized, anatomically accurate spiking network simulator," in *Biomedical Sciences and Engineering Conference (BSEC), 2010* (Oak Ridge, TN: IEEE), 1–3.

Takaya, Y. U., and Arita, T. (2003). "Situated and embodied evolution in collective evolutionary robotics," in *Proc. of the 8th International Symposium on Artificial Life and Robotics* (Beppu: Citeseer).

Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics. Intelligent Robotics and Autonomous Agents*. Cambridge, MA: MIT Press.

Trianni, V., Groß, R., Labella, T. H., Sahin, E., and Dorigo, M. (2003). "Evolving aggregation behaviors in a swarm of robots," in *Advances in Artificial Life* (Berlin: Springer), 865–874.

Trianni, V., and Nolfi, S. (2011). Engineering the evolution of self-organizing behaviors in swarm robotics: a case study. *Artif. Life* 17, 183–202. doi:10.1162/artl_a_00031

Vanderelst, D., and Winfield, A. (2018). An architecture for ethical robots inspired by the simulation theory of cognition. *Cogn. Syst. Res.* 48, 56–66. doi:10.1016/j.cogsys.2017.04.002

Vásárhelyi, G., Virágh, C., Somorjai, G., Tarcai, N., Szörényi, T., Nepusz, T., et al. (2014). "Outdoor flocking and formation flight with autonomous aerial robots," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on* (Chicago, IL: IEEE), 3866–3873.

Vaughan, R. (2008). Massively multi-robot simulation in stage. *Swarm Intell.* 2, 189–208. doi:10.1007/s11721-008-0014-4

Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: distributing an evolutionary algorithm in a population of robots. *Rob. Auton. Syst.* 39, 1–18. doi:10.1016/S0921-8890(02)00170-7

White, J. G., Southgate, E., Thomson, J. N., and Brenner, S. (1986). The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 314, 1–340. doi:10.1098/rstb.1986.0056

Whitley, D., Rana, S., and Heckendorn, R. B. (1999). The island model genetic algorithm: on separability, population size and convergence. *CIT J. Comput. Inform. Technol.* 7, 33–47.

Williams, R. L., Carter, B. E., Gallina, P., and Rosati, G. (2002). Dynamic model with slip for wheeled omnidirectional robots. *IEEE Trans. Robot. Autom.* 18, 285–293. doi:10.1109/TRA.2002.1019459

Wilson, S., Gameros, R., Sheely, M., Lin, M., Dover, K., Gevorkyan, R., et al. (2016). Pheeno, a versatile swarm robotic research and education platform. *IEEE Robot. Autom. Lett.* 1, 884–891. doi:10.1109/LRA.2016.2524987

Winfield, A. F. (2015). "Robots with internal models: a route to self-aware and hence safer robots," in *The Computer after Me: Awareness and Self-Awareness in Autonomic Systems* (London, UK: World Scientific), 237–252.

Winfield, A. F., Blum, C., and Liu, W. (2014). "Towards an ethical robot: internal models, consequences and ethical action selection," in *TAROS 2014 – Towards Autonomous Robotic Systems*, volume 8717 of *Lecture Notes in Computer Science*, eds M. Mistry, A. Leonardis, M. Witkowski, and C. Melhuish (Berlin: Springer International Publishing), 85–96.

Xianyi, Z., Qian, W., and Yunquan, Z. (2012). "Model-driven level 3 BLAS performance optimization on Loongson 3A processor," in *Parallel and Distributed Systems (ICPADS), 2012 IEEE 18th International Conference on* (Singapore: IEEE), 684–691.

Zagal, J. C., Ruiz-del Solar, J., and Vallejos, P. (2004). "Back to reality: crossing the reality gap in evolutionary robotics," in *IAV 2004 the 5th IFAC Symposium on Intelligent Autonomous Vehicles* (Lisbon, Portugal).

frontiers
in Robotics and AI

Check for
updates

Embodied Evolution in Collective Robotics: A Review

Nicolas Bredeche[1]*, Evert Haasdijk[2] and Abraham Prieto[3]

[1] *Sorbonne Université, CNRS, Institute of Intelligent Systems and Robotics, ISIR, Paris, France,* [2] *Computational Intelligence Group, Department of Computer Science, Vrije Universiteit, Amsterdam, Netherlands,* [3] *Integrated Group for Engineering Research, Universidade da Coruña, Ferrol, Spain*

This article provides an overview of evolutionary robotics techniques applied to online distributed evolution for robot collectives, namely, embodied evolution. It provides a definition of embodied evolution as well as a thorough description of the underlying concepts and mechanisms. This article also presents a comprehensive summary of research published in the field since its inception around the year 2000, providing various perspectives to identify the major trends. In particular, we identify a shift from considering embodied evolution as a parallel search method within small robot collectives (fewer than 10 robots) to embodied evolution as an online distributed learning method for designing collective behaviors in swarm-like collectives. This article concludes with a discussion of applications and open questions, providing a milestone for past and an inspiration for future research.

Keywords: embodied evolution, online distributed evolution, collective robotics, evolutionary robotics, collective adaptive systems

1. INTRODUCTION

This article provides an overview of evolutionary robotics research where evolution takes place in a population of robots in a continuous manner. Ficici et al. (1999) coined the phrase *embodied evolution* for evolutionary processes that are distributed over the robots in the population to allow them to adapt autonomously and continuously. As robotics technology becomes simultaneously more capable and economically viable, individual robots operated at large expense by teams of experts are increasingly supplemented by collectives of robots used cooperatively under minimal human supervision (Bellingham and Rajan, 2007), and embodied evolution can play a crucial role in enabling autonomous online adaptivity in such robot collectives.

The vision behind embodied evolution is one of the collectives of truly autonomous robots that can adapt their behavior to suit varying tasks and circumstances. Autonomy occurs at two levels: not only the robots perform their tasks without external control but also they assess and adapt—through evolution—their behavior without referral to external oversight and so learn autonomously. This adaptive capability allows robots to be deployed in situations that cannot be accurately modeled *a priori*. This may be because the environment or user requirements are not fully known, or it may be due to the complexity of the interactions among the robots as well as with their environment effectively rendering the scenario unpredictable. Also, onboard adaptivity intrinsically avoids the reality gap (Jakobi et al., 1995) that results from inaccurate modeling of robots or their environment when developing controllers before deployment because controllers continue to develop *after* deployment. A final benefit is that embodied evolution can be seen as parallelizing the evolutionary process because it distributes the evaluations over multiple robots. Alba (2002) has shown that such parallelism can provide substantial benefits, including superlinear speedups. In the case of robots,

this has the added benefit of reducing the amount of time spent executing poor controllers per robot, reducing wear and tear.

Embodied evolution's online nature contrasts with "traditional" evolutionary robotics research. Traditional evolutionary robotics employs evolution in the classical sequential centralized optimization paradigm: parent and survivor selection are centralized and consider the entire population. The "robotics" part entails a series of robotic trials (simulated or not) in an evolution-based search for optimal robot controllers (Nolfi and Floreano, 2000; Bongard, 2013; Doncieux et al., 2015). In terms of task performance, embodied evolution has been shown to outperform alternative evolutionary robotic techniques in some setups such as surveillance and self-localization with flying UAVs (Schut et al., 2009; Prieto et al., 2016), especially regarding convergence speed.

To provide a basis for a clear discussion, we define embodied evolution as a paradigm where evolution is implemented in multirobotic (two or more robots) system. Two robots are already considered a multirobotic system since it is still possible to distribute an algorithm among them. These systems exhibit the following features.

Decentralized: There is no central authority that selects parents to produce offspring or individuals to be replaced. Instead, robots assess their performance, exchange, and select genetic material autonomously on the basis of locally available information.

Online: Robot controllers change on the fly, as the robots go about their proper actions: evolution occurs during the operational lifetime of the robots and in the robots' task environment. The process continues after the robots have been deployed.

Parallel: Whether they collaborate in their tasks or not, the population consists of multiple robots that perform their actions and evolve concurrently, in the same environment, interacting frequently to exchange genetic material.

The decentralized nature of communicating genetic material implies that the selection is executed locally, usually involving only a part of the whole population (Eiben et al., 2007), and that it must be performed by the robots themselves. This adds a third opportunity for selection in addition to parent and survivor selection as defined for classical evolutionary computing. Thus, embodied evolution extends the collection of operators that define an evolutionary algorithm (i.e., evaluation, selection, variation, and replacement (Eiben and Smith, 2008)) with *mating* as a key evolutionary operator.

Mating: An action where two (or more) robots decide to send and/or receive genetic material, whether this material will or will not be used for generating new offspring. When and how this happens depends not only on predefined heuristics but also on evolved behavior, the latter determining to a large extent whether robots ever meet to have the opportunity to exchange genetic material.

In the past 20 years, online evolutionary robotics in general and embodied evolution in particular have matured as research fields. This is evidenced by the growing number of relevant publications in respected evolutionary computing venues such as in conferences (e.g., ACM GECCO, ALIFE, ECAL, and EvoApplications), journals (e.g., Evolutionary Intelligence's special issue on Evolutionary Robotics (Haasdijk et al., 2014b)), workshops (PPSN 2014 ER workshop, GECCO 2015 and 2017

Evolving collective behaviors in robotics workshop), and tutorials (ALIFE 2014, GECCO 2015 and 2017, ECAL 2015, PPSN 2016, and ICDL-EPIROB 2016). A Google Scholar search of publications citing the seminal embodied evolution paper by Watson et al. (2002) illustrates this growing trend. Since 2009, the paper has attracted substantial interest, more than doubling the yearly number of citations since 2008 (approximately 20 citations per year since then).[1]

To date, however, a clear definition of what embodied evolution is (and what it is not) and an overview of the state of the art in this area are not available. This article provides a definition of the embodied evolution paradigm and relates it to other evolutionary and swarm robotics research (Sections 2 and 3). We identify and review relevant research, highlighting many design choices and issues that are particular to the embodied evolution paradigm (Sections 4 and 5). Together this provides a thorough overview of the relevant state-of-the-art and a starting point for researchers interested in evolutionary methods for collective autonomous adaptation. Section 6 identifies open issues and research in other fields that may provide solutions, suggests directions for future work, and discusses potential applications.

2. CONTEXT

Embodied evolution considers collectives of robots that adapt online. This section positions embodied evolution vis à vis other methods for developing controllers for robot collectives and for achieving online adaptation.

2.1. Offline Design of Behaviors in Collective Robotics

Decentralized decision-making is a central theme in collective robotics research: when the robot collective cannot be centrally controlled, the individual robots' behavior must be carefully designed so that global coordination occurs through local interactions.

Seminal works from the 1990s such as Mataric's Nerd Herd (Mataric, 1994) addressed this problem by hand-crafting behavior-based control architectures. Manually designing robot behaviors has since been extended with elaborate methodologies and architectures for multirobot control (see Parker (2008) for a review) and with a plethora of bioinspired control rules for swarm-like collective robotics (see Nouyan et al. (2009) and Rubenstein et al. (2014) for recent examples involving real robots and Beni (2005), Brambilla et al. (2012), and Bayindir (2016) for discussions and recent reviews).

Automated design methods have been explored with the hope of tackling problems of greater complexity. Early examples of this approach were applied to the robocup challenge for learning coordination strategies in a well-defined setting. See the study by Stone and Veloso (1998) for an early review and Stone et al. (2005) and Barrett et al. (2016) for more recent work in this vein. However, Bernstein et al. (2002) demonstrated that solving even the simplest multiagent learning problem is intractable in

---

[1]See https://plot.ly/~evertwh/17/ for more details and the underlying data.

polynomial time (actually, it is NEXP-complete), so obtaining an optimal solution in reasonable time is currently infeasible. Recent works in reinforcement learning have developed theoretical tools to break down complexity by operating a move from considering many agents to a collection of single agents, each of which being optimized separately (Dibangoye et al., 2015), leading to theoretically well-founded contributions, but with limited practical validation involving very few robots and simple tasks (Amato et al., 2015).

Lacking theoretical foundations, but instead based on the experimental validation, swarm robotics controllers have been developed with black-box optimization methods ranging from brute-force optimization using a simplified (hence tractable) representation of a problem (Werfel et al., 2014) and evolutionary robotics (Hauert et al., 2008; Trianni et al., 2008; Gauci et al., 2012; Silva et al., 2016).

The methods vary, but all the approaches described here (including "standard" evolutionary robotics) share a common goal: to design or optimize a set of control rules for autonomous robots that are part of a collective *before* the actual deployment of the robots. The particular challenge in this kind of work is to design individual behaviors that lead to some required global ("emergent") behavior without the need for central oversight.

2.2. Lifelong Learning in Evolutionary Robotics

It has long been argued that deploying robots in the real world may benefit from continuing to acquire new capabilities *after* initial deployment (Thrun and Mitchell, 1995; Nelson and Grant, 2006), especially if the environment is not known beforehand. Therefore, the question we are concerned with in this article is *how to endow a collective robotics system with the capability to perform lifelong learning*. Evolutionary robotics research into this question typically focuses on individual autonomous robots. Early works in evolutionary robotics that considered lifelong learning explored learning mechanisms to cope with minor environmental changes (see the classic book by Nolfi and Floreano (2000) and Urzelai and Floreano (2001) and (Tonelli and Mouret, 2013) for examples and Mouret and Tonelli (2015) for a nomenclature). More recently, Bongard et al. (2006) and Cully et al. (2015) addressed

resilience by introducing fast online re-optimization to recover from hardware damage.

Bredeche et al. (2009), Christensen et al. (2010), and Silva et al. (2012) are some examples of online versions of evolutionary robotics algorithms that target the fully autonomous acquisition of behavior to achieve some predefined task in individual robots. Targeting agents in a video game rather than robots, Stanley et al. (2005) tackled the online evolution of controllers in a multiagent system. Because the agents were virtual, the researchers could control some aspects of the evaluation conditions (e.g., restarting the evaluation of agents from the same initial position). This kind of control is typically not feasible in autonomously deployed robotic systems.

Embodied evolution builds on evolutionary robotics to implement lifelong learning in robot *collectives*. Its clear link with traditional evolutionary robotics is exemplified by work such as by Usui and Arita (2003), where a traditional evolutionary algorithm is encapsulated on each robot. Individual controllers are evaluated sequentially in a standard time sharing setup, and the robots implement a communication scheme that resembles an island model to exchange genomes from one robot to another. It is this communication that makes this an instance of embodied evolution.

3. ALGORITHMIC DESCRIPTION

This section presents a formal description of the embodied evolution paradigm by means of generic pseudocode and a discussion about its operation from a more conceptual perspective.

The pseudocode in **Algorithm 1** provides an idealized description of a robot's control loop as it pertains to embodied evolution. Each robot runs its own instance of the algorithm, and the evolutionary process emerges from the interaction between the robots. In embodied evolution, there is no entity outside the robots that oversees the evolutionary process, and there is typically no synchronization between the robots: the replacement of genomes is asynchronous and autonomous.

Some steps in this generic control loop can be implicit or entwined in particular implementations. For instance, robots may continually broadcast genetic material over short range, so that other robots that come within this range receive it

---

ALGORITHM 1 | An individual robot's control loop for embodied evolution.

```
initialize robot;
for ever do
    Sense - Act cycle (depends on robotic paradigm);
    perf ← calculate performance;

    if mating? then                                       //E.g., is another robot nearby?
        transmit my genome;                               //and optional further information
        g ← receive mate's genome;
        store(g);
    end

    if replacement? then                                  //E.g., time or virtual energy runs out
        parents ← select parents;
        offspring ← variation(parents);
        activate(offspring)//Time-sharing: control is handed over to the new candidate controller
    end
end
```

automatically. In such a case, the *mating* operation is implicitly defined by the selected broadcast range. Similarly, genetic material may be incorporated into the currently active genome as it is received, merging the mating and replacement operations. Implicitly defined or otherwise, the steps in this algorithm are, with the possible exception of performance calculation, necessary components of any embodied evolution implementation.

The following list describes and discusses the steps in the algorithm in detail.

Initialization: The robot controllers are typically initialized randomly, but it is possible that the initial controllers are developed offline, be it through evolution or handcraft (e.g., see the work by Hettiarachchi et al. (2006)).

Sense–act cycle: This represents "regular," i.e., not related to the evolutionary process, robot control. The details of the sense–act cycle depend on the robotic paradigm that governs robot behavior; this may include planning, subsumption, or other paradigms. This may also be implemented as a separate parallel process.

Calculate performance: If the evolutionary process defines an objective function, the robots monitor their own performance. This may involve measurements of quantities such as speed, number of collisions, or amount of collected resources. Whatever their nature, these measurements are then used to evaluate and compare genomes (as fitness values in evolutionary computation). The possible discrepancy between the individual's objective function and the population welfare will be discussed further in Section 6.2.

Mating: This is the essential step in the evolutionary process where robots exchange genetic material. The choice to mate with another robot may be purely based on the environmental contingencies (e.g., when robots mate whenever they are within communication range), but other considerations may also play a part (e.g., performance, genotypic similarity). The pseudocode describes a symmetric exchange of genomes (both with a transmit and a receive operation), but this may be asymmetrical for particular implementations. In implementations such as that of Schwarzer et al. (2011) or Haasdijk et al. (2014a), for instance, robots suspend normal operation to collect genetic material from other, active robots. Mating typically results in a pool of candidate parents that are considered in the parent selection process.

Replacement: The currently active genome is replaced by a new individual (the offspring), implying the removal of the current genome. This event can be triggered by a robot's internal conditions (e.g., running out of time or virtual energy, reaching a given performance level) or through interactions with other robots (e.g., receiving promising genetic material (Watson et al., 2002)).

Parent selection: This is the process that selects which genetic information will be used for the creation of new offspring from the received genetic information through mating events. When an objective is defined, the performance of the received genome is usually the basis for selection, just as in regular evolutionary computing. In other cases, the selection among received genomes can be random or depend on non-performance related heuristics (e.g., random, genotypic proximity). In the absence of objective-driven selection pressure, individuals are still competing with

respect to their ability to spread their own genome within the population, although that cannot be explicitly captured during parent selection. This will be further discussed in Section 5.2.

Variation: A new genome is created by applying the variation operators (mutation and crossover) on the selected parent genome(s). This is subsequently activated to replace the current controller.

From a conceptual perspective, embodied evolution can be analyzed at two levels that are represented by two as depicted in **Figure 1**.

The robot-centric cycle is depicted on the right in **Figure 1**. It represents the physical interactions that occur between the robot and its environment, including interactions with other robots and extends this sense–act loop commonly used to describe real-time control systems by accommodating the exchange and activation of genetic material. At this particular point, the genome-centric and robot-centric cycles overlap. The cycle operates as follows: each robot is associated with an *active* genome, and the genome is interpreted into a set of features and control architecture (the phenotype), which produces a behavior that includes the transmission of its own genome to some other robots. Each robot eventually switches from an active genome to another, depending on a specific event (e.g., minimum energy threshold) or duration (e.g., fixed lifetime), and consequently changes its active genome, probably impacting its behavior.

The genome-centric cycle deals with the events that directly affect the genomes existing in the robot population and therefore also the evolution *per se*. Again, the mating and the replacement are the events that overlap with the robot-centric cycle. The operation from the genome cycle perspective is as follows: each robot starts with an initial genome, either initialized randomly or *a priori* defined. While this genome is *active*, it determines the phenotype of the robot, hence its behavior. Afterward, when replacement is triggered, some genomes are selected from the reservoir of genomes previously received according to the parent selection criteria and later combined using the variation operators. This new genome will then become part of the population. In the case of fixed-size population algorithms, the replacement will automatically trigger the removal of the old genome. In some other cases, however, there is a specific criterion to trigger the removal event producing populations of individuals that change their size along the evolution.

The two circles connect on two occasions, first by the "exchange genomes" (or mating) process, which implies the transmission of genetic material, possibly together with additional information (fitness if available, general performance, genetic affinity, etc.) to modulate future selection. Generally, the received information is stored to be used (in full or in part) to replace the active genome in the later parent selection process. Therefore, the event is triggered and modulated by the robot cycle, but it impacts on the genomic cycle. Also, the decentralized nature of the paradigm enforces that these transmissions occur locally, either one-to-one or to any robot in a limited range. There are several ways in which mate selection can be implemented, for instance, individuals may send and receive genomic information indiscriminately within a certain location range or the frequency of transmission can depend on the task performance. The second overlap between
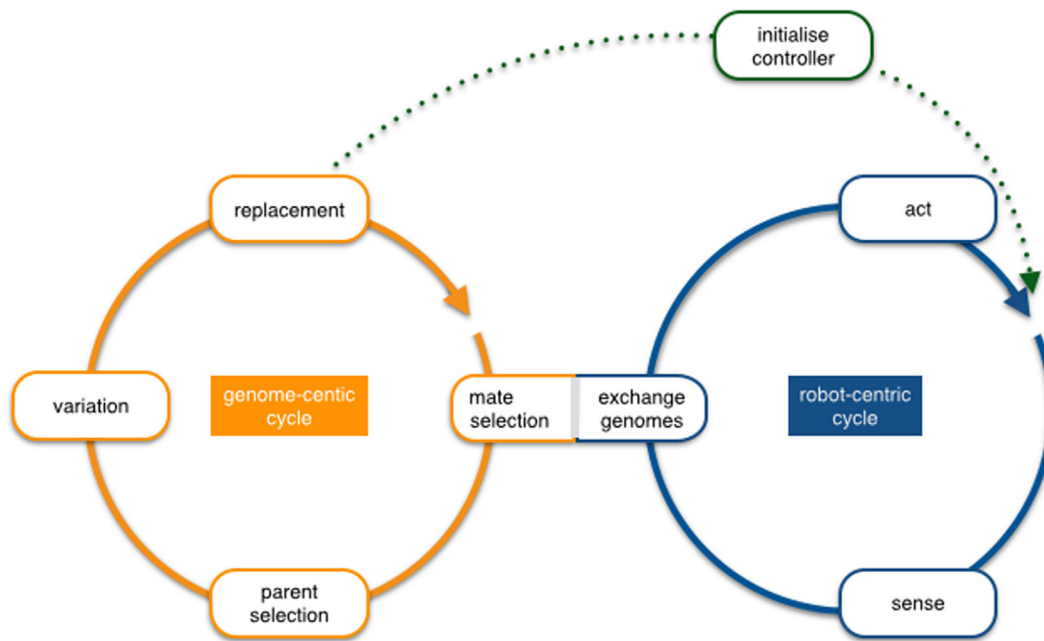
FIGURE 1 | The overlapping robot-centric and genome-centric cycles in embodied evolution. The robot-centric cycle uses a single *active* genome that determines the current robot behavior (sense–act loop), and the genome-centric cycle manages an internal *reservoir* of genomes received from other robots or built locally (parent selection/variation), of which the next active genome will be selected eventually (replacement).

the two cycles is the activation of new genomic information (replacement). The activation of a genome in the genomic cycle implies that this new genome will now take control of the robot and therefore changes the response of the robot in the scenario (in evolutionary computing terms, this event will mark the start of a new individual evaluation). This aspect is what creates the online character of the algorithm that, together with the locality constraints, implies that the process is also asynchronous.

This conceptual representation matches what has been defined as *distributed* embodied evolution by Eiben et al. (2010). The authors proposed a taxonomy for online evolution that differentiates between encapsulated, distributed, and hybrid schemes. Most embodied evolution implementations are distributed, but this schematic representation also covers hybrid implementations. In such cases, the robot locally maintains a population that is augmented through mating (rather like an island model in parallel evolutionary algorithms). It should be noted that encapsulated implementations (where each robot runs independently of the others) are not considered in this overview.

4. EMBODIED EVOLUTION: THE STATE OF THE ART

In this section, we identify the major research topics from the works published since the inception of the domain, all summarized in **Table 1**. **Table 1** provides an overview of published research on embodied evolution with robot collectives. Each entry describes a contribution, which may cover several papers. The entries are described in terms of their implementation details, the robot behavior, experimental settings, mating conditions, selection, and replacement schemes. The glossary in **Table 2** explains these features in more detail.

First, we distinguish between works that consider embodied evolution as a parallel search method for optimizing *individual* behaviors and works where embodied evolution is employed to craft *collective* behavior in robot populations. The latter trend, where the emphasis is on collective behavior, has emerged relatively recently and since then has gained importance (32 papers since 2009).

Second, we consider the homogeneity of the evolving population; borrowing definitions from biology, we use the term *monomorphic* (resp. *polymorphic*) for a population containing one (resp. more than one) class of genotype, for instance, to achieve specialization. A monomorphic population implies that individuals will behave in a similar manner (except for small variations due to minor genetic differences). On the contrary, polymorphic populations host multiple groups of individuals, each group with its particular genotypic signature, possibly displaying a specific behavior. Research to date shows that cooperation in monomorphic populations can be easily achieved (e.g. (Prieto et al., 2010; Schwarzer et al., 2010; Montanier and Bredeche, 2011, 2013; Silva et al., 2012)), while polymorphic populations (e.g., displaying genetic-encoded behavioral specialization) require very specific conditions to evolve (e.g., Trueba et al. (2013); Haasdijk et al. (2014a); Montanier et al. (2016)).

A notable number of contributions employ real robots. Since the first experiments in this field, the intrinsic online nature of embodied evolution has made such validation comparatively straightforward (Ficici et al., 1999; Watson et al., 2002). "Traditional" evolutionary robotics is more concerned with

TABLE 1 | Overview of Embodied Evolution research.

	Implementation				Robot behavior					Experimental settings		Mating conditions			Selection scheme			Replacement scheme			
	Distributed	Hybrid	Simulation	Real robots	Monomorphic	Polymorphic	Individual behavior	Cooperation	Division of labor	Task	No of robots	Panmictic	Proximity	Other	Performance	Random	Genotypic distance	Fixed lifetime	Variable lifetime	Event based	Limited lifetime
Ficici et al. (1999); Watson et al. (2002)	●		●	●			●			Phototaxis	8		●		●				●		
Simões and Dimond (2001)	●		●	●			●			Obstacle avoidance	6	●			●			●			
Usui and Arita (2003)		●	●	●	●		●			obstacle avoidance	6		●		●			●			
Bianco and Nolfi (2004)	●		●	●			●			Self-assembly	64		●			●				●	
Hettiarachchi et al. (2006); Hettiarachchi and Spears (2009)	●		●	●			●			Navigation with obstacle avoidance	60		●		●					●	
Wischmann et al. (2007)	●		●	●			●			Foraging[a]	3		●		●				●		
Perez et al. (2008)		●	●	●			●			Obstacle avoidance	5	●			●			●			
König and Schmeck (2009); König et al. (2009)	●		●	●			●			Obstacle avoidance with gate passing	26; 30		●		●			●			
Pugh and Martinoli (2009)		●	●	●	●		●			Obstacle avoidance	1–10	●	●		●			●			
Prieto et al. (2009); Trueba et al. (2011, 2012)	●		●			●		●	●	Surveillance, foraging, construction	20		●		●					●	●
Bredeche and Montanier (2010, 2012) Bredeche (2014)	●		●	●			●			None	20; 4000		●			●					●
Prieto et al. (2010)	●		●		●			●	●	Surveillance	8		●		●					●	●
Schwarzer et al. (2010)	●		●	●	●		●			None	Up to 40		●			●			●		
Schwarzer et al. (2011)	●		●		●	●				Phototaxis	4		●		●			●			
Montanier and Bredeche (2011, 2013)	●		●	●				●		None	100		●			●	●				●
Huijsman et al. (2011)	●	●	●		●		●			Obstacle avoidance	4–400	●		●	●			●			

(*Continued*)

TABLE 1 | Continued

	Implementation				Robot behavior					Experimental settings		Mating conditions			Selection scheme			Replacement scheme			
	Distributed	Hybrid	Simulation	Real robots	Monomorphic	Polymorphic	Individual behavior	Cooperation	Division of labor	Task	No of robots	Panmictic	Proximity	Other	Performance	Random	Genotypic distance	Fixed lifetime	Variable lifetime	Event based	Limited lifetime
Karafotias et al. (2011)		•	•		•	•			•	Obstacle avoidance, phototaxis, and patrolling	10		•		•			•			
Silva et al. (2012, 2013, 2015, 2017)		•	•	•	•	•		•		Navigation, aggregation, surveillance, and phototaxis	2–20	•			•					•	
Weel et al. (2012a,b)		•	•		•			•		Foraging	10; 50	•			•			•			
García-Sánchez et al. (2012)		•	•		•		•			Obstacle avoidance	4–36	•	•		•			•			
Haasdijk and Bredeche (2013) Haasdijk et al. (2013, 2014a); Noskov et al. (2013); Haasdijk and Eigenhuis (2016); Bangel and Haasdijk (2017); Kemeling and Haasdijk (2017)	•		•			•	•		•	Foraging	100	•			•			•			
Trueba et al. (2013) Trueba (2017)	•		•	•	•			•	•	Synthetic mapping, gathering, self-location	40; 20; 9	•	•		•	•				•	•
O'Dowd et al. (2014)		•	•	•	•		•			Foraging	10	•			•			•			
Fernandez Pérez et al. (2014)	•		•	•					•	Foraging	50	•			•	•		•			
Fernandez Pérez et al. (2015)	•		•	•				•		Foraging	100	•			•					•	
Hart et al. (2015); Steyven et al. (2016)	•		•	•				•		Foraging	100	•			•						•
Heinerman et al. (2015, 2016)		•	•	•	•		•			Obstacle avoidance	6	•			•			•			
Montanier et al. (2016); Bredeche et al. (2017)	•		•				•		•	Foraging	100; 500	•			•	•					•
Fernandez Pérez et al. (2017)	•		•				•	•		Foraging	200	•			•			•			

aAs a proxy for predator avoidance.

TABLE 2 | Glossary.

Field	Comment
Implementation	**Distributed** implementations have one genome for each robot, and an offspring is created only as the result of a mating event or by mutating the current genome. **Hybrid** implementations have multiple genomes per robot, and offspring can be created from this internal pool and from genomes "imported" through mating events. As stated earlier, the encapsulated scheme is not considered embodied evolution as there is no exchange of genomes between robots in this case.
	The experiments can use **real robots** or **simulation**.
Robot behavior	A **monomorphic** population contains individuals with similar genotypes (with variations due to mutation). A **polymorphic** population is divided into two (or more) subgroups of genetically similar individuals, and different genotypic signatures from one group to the other, e.g., to achieve specialization.
	We distinguish between experiments that target efficient **individual behavior** vs. collective behaviors (i.e., social behaviors, incl. cooperation)
Experimental settings	Identifies the **task(s)** considered in the experiment, e.g., obstacle avoidance, foraging, … **None** indicates that there is no user-defined task and that consequently, selection pressure results from the environment only. The **number of robots** used is also included. $n_1 - n_2$ indicates the interval for one experiment and $n1$; $n2$ gives numbers for two experiments.
Mating conditions	Mating can be based on **proximity**: two robots can mate whenever they are physically close to each other (e.g., in infrared communication range). In **panmictic** systems, robots can mate with all other robots, regardless of their location. **Other** comprises systems where robots maintain an explicit list of potential mates (a social network), which may be maintained through gossiping.
Selection scheme	Parents are selected from the received and internal genomes on the basis of their **performance** if a task is defined. **Random** parent selection implies only environment-driven selection. Currently, the only examples of other selection schemes use genotypic distance, but this category also covers metrics such as novelty.
Replacement scheme	Genomes can have a **fixed lifetime**, **variable lifetime**, or **limited lifetime** (similar to variable lifetime, but with an upper bound). **Event-based** replacement schemes do not depend on time but on events such as reception of genetic material (e.g., in the microbial GA used by Watson et al. (2002)).

robustness at the level of the evolved *behavior* (mostly caused from the reality gap that exists between simulation and the real world) than is embodied evolution, which emphasizes the design of robust *algorithms*, where transfer between simulation and real world may be less problematic. In the contributions presented here, simulation is used for extensive analysis that could hardly take place with real robots due to time or economic constraint. Still, it is important to note that many researchers who use simulation have also published works

with real robots, thus including real-world validation in their research methodology.

Since 2010, there have been a number of experiments that employ large ($\geq$100) numbers of (simulated) robots, shifting toward more swarm-like robotics where evolutionary dynamics can be quite different (Huijsman et al., 2011; Bredeche, 2014; Haasdijk et al., 2014b). Recent works in this vein focus on the nature of selection pressure, emphasizing the unique aspect of embodied evolution that selection pressure results from both the environment (which impacts mating) and the task. Bredeche and Montanier (2010, 2012) showed that environmental pressure alone can drive evolution toward self-sustaining behaviors. Haasdijk et al. (2014a) showed that these selection pressures can to some extent be modulated by tuning the ease with which robots can transmit genomes. Steyven et al. (2016) showed that adjusting the availability and value of energy resources results in the evolution of a range of different behaviors. These results emphasize that tailoring the environmental requirements to transmit genomes can profoundly impact the evolutionary dynamics and that understanding these effects is vital to effectively develop embodied evolution systems.

5. ISSUES IN EMBODIED EVOLUTION

What sets embodied evolution apart from classical evolutionary robotics (and, indeed, from most evolutionary computing) is the fact that evolution acts as a force for continuous adaptation, not (just) as an optimizer before deployment. As a continuous evolutionary process, embodied evolution is similar to some evolutionary systems considered in artificial life research (e.g. Axelrod (1984); Ray (1993), to name a few). The operations that implement the evolutionary process to adapt the robots' controllers are an integral part of their behavior in their task environment. This includes mating behavior to exchange and select genetic material, assessing one's own and/or each other's task performance (if a task is defined) and applying variation operators such as mutation and recombination.

This raises issues that are particular to embodied evolution. The research listed in the previous section has identified and investigated a number of these issues, and the remainder of this section discusses these issues in detail, while Section 6.2 discusses issues that so far have not benefited from close attention in embodied evolution research.

5.1. Local Selection
In embodied evolution, the evolutionary process is generally implemented through local interactions between the robots, i.e., the mating operation introduced above. This implies the concept of a neighborhood from which mates are selected. One common way to define neighborhood is to consider robots within communication range, but it can also be defined in terms of other distance measures such as genotypic or phenotypic distance. Mates are selected by sampling from this neighborhood, and a new individual is created by applying variation operators to the sampled genome(s). This local interaction has its origin in constraints that derive from communication limitations in some distributed robotic scenarios. Schut et al. (2009) showed it to

be beneficial in simulated setups as an exploration/exploitation balancing mechanism.

Embodied evolution, with chance encounters providing the sampling mechanism, has some similarities with other flavors of evolutionary computation. Cellular evolutionary algorithms (Alba and Dorronsoro, 2008) consider continuous random rewiring of a network topology (in a grid of CPUs or computers) where all elements are evaluated in parallel. In this context, locally selecting candidates for reproduction is a recurring theme that is shared with embodied evolution (e.g., García-Sánchez et al. (2012); Fernandez Pérez et al. (2014)).

5.2. Objective Functions vs Selection Pressure

In traditional evolutionary algorithms, the optimization process is guided by a (set of) objective function(s) (Eiben and Smith, 2008). Evaluation of the candidate solutions, i.e., of the genomes in the population, allows for (typically numerical) comparison of their performance. Beyond its relevance for performance assessment, the evaluation process *per se* has generally no influence on the manner in which selection, variation, and replacement evolutionary operators are applied. This is different in embodied evolution, where the behavior of an individual can directly impact the likelihood of encounters with others and so influence selection and reproductive success (Bredeche and Montanier, 2010). Evolution can not only improve task performance but can also develop mating strategies, for example, by maximizing the number of encounters between robots if that improves the likelihood of transmitting genetic material.

It is therefore important to realize that the *selection pressure* on the robot population does not only derive from the specified *objective function(s)* as it traditionally does in evolutionary computation. In embodied evolution, the environment, including the mechanisms that allow mating, also exert selection pressure. Consequently, evolution experiences selection pressure from the aggregate of objective function(s) and environmental particularities. Steyven et al. (2016) researched how aspects of the robots' environment influence the emergence of particular behaviors and the balance between pressure toward survival and task. The objective may even pose requirements that are opposed to those by the environment. This can be the case when a task implies risky behaviors or because a task requires resources that are also needed for survival and mating. In such situations, the evolutionary process must establish a tradeoff between objective-driven optimization and the maintenance of a viable environment where evolution occurs, which is a challenge in itself (Haasdijk, 2015).

5.3. Autonomous Performance Evaluation

The decentralized nature of the evolutionary process implies that there is no omniscient presence who knows (let alone determines) the fitness values of all individuals. Consequently, when an objective function is defined, it is the robots themselves that must gage their performance and share it with other robots when mating: each robot must have an evaluation function that can be computed onboard and autonomously. Typical examples of such evaluation functions are the number of resources collected, the number of times a target has been reached, or the number of collisions. The requirement of autonomous assessment does not fundamentally change the way one defines fitness functions, but it does impact their usage as shown by Nordin and Banzhaf (1997), Walker et al. (2006), and Wolpert and Tumer (2008).

Evaluation time: The robots must run a controller for some time to assess the resultant behavior. This implies a *time sharing* scheme where robots run their current controllers to evaluate their performance. In many similar implementations, a robot runs a controller for a fixed evaluation time; Haasdijk et al. (2012) showed that this is a very important parameter in encapsulated online evolution, and it is likely to be similarly influential in embodied evolution.

Evaluation in varying circumstances: Because the evolutionary machinery (mating, evaluating new individuals, etc.) is an integral part of robot behavior, which runs in parallel with the performance of regular tasks, there can be no thorough re-initialization or re-positioning procedure between genome replacements. This implies a noisy evaluation: a robot may undervalue a genome starting in adverse circumstances and vice versa. As Nordin and Banzhaf (1997) (p. 121) put it: "*Each individual is tested against a different real-time situation leading to a unique fitness case. This results in 'unfair' comparison where individuals have to navigate in situations with very different possible outcomes. However, […] experiments show that over time averaging tendencies of this learning method will even out the random effects of probabilistic sampling and a set of good solutions will survive.*" Bredeche et al. (2009) proposed a re-evaluation scheme to address this issue: seemingly efficient candidate solutions have a probability to be re-evaluated to cope with possible evaluation noise. A solution with a higher score *and* a lower variance will then be preferred to one with a higher variance. While re-evaluation is not always used in embodied evolution, the evaluation of relatively similar genomes on different robots running in parallel provides another way to smooth the effect of noisy evaluations.

Multiple objectives: To deal with multiple objectives, evolutionary computation techniques typically select individuals on the basis of Pareto dominance. While this is eminently possible when selecting partners as well, Pareto dominance can only be determined vis à vis the population sample that the selecting robot has acquired. It is unclear how this affects the overall performance and if the robot collective can effectively cover the Pareto front. Bangel and Haasdijk (2017) investigated the use of a "market mechanism" to balance the selection pressure over multiple tasks in a concurrent foraging scenario, showing that this at least prevents the robot collective from focusing on single tasks, but that it does not lead to specialization in individual robots.

6. DISCUSSION

The previous sections show that there is a considerable and increasing amount of research into embodied evolution, addressing issues that are particular to its autonomous and distributed nature. This section turns to the future of embodied evolution research, discussing potential applications and proposing a

research agenda to tackle some of the more relevant and immediate issues that so far have remained insufficiently addressed in the field.

6.1. Applications of Embodied Evolution

Embodied evolution can be used as a design method for engineering, as a modeling method for evolutionary biology, or as a method to investigate evolving complex systems more generally. Let us briefly consider each of these possibilities.

Engineering: The online adaptivity afforded by embodied evolution offers many novel possibilities for deployment of robot collectives: exploration of unknown environments, search and rescue, distributed monitoring of large objects or areas, distributed construction, distributed mining, etc. Embodied evolution can offer a solution when robot collectives are required to be versatile, since the robots can be deployed in and adapt to open and *a priori* unknown environments and tasks. The collective is comparatively robust to failure through redundancy and the decentralized nature of the algorithm because the system continues to function even if some robots break down. Embodied evolution can increase autonomy because the robots can, for instance, learn how to maintain energy while performing their task without intervention by an operator.

Currently, embodied evolution has already provided solutions to tasks such as navigation, surveillance, and foraging (see **Table 1** for a complete list), but these are of limited interest because of the simplicity of the tasks considered in research to date. The research agenda proposed in Section 6.2 provides some suggestions for further research to mitigate this.

Evolutionary biology: In the past 100 years, evolutionary biology benefited from both experimental and theoretical advances. It is now possible, for instance, to study evolutionary mechanisms through methods such as gene sequencing (Blount et al., 2012; Wiser et al., 2013). However, *in vitro* experimental evolution has its limitations: with evolution in "real" substrates, the time scales involved limit the applicability to relatively simple organisms such as *Escherichia coli* (Good et al., 2017). From a theoretical point of view, population genetics (see Charlesworth and Charlesworth (2010) for a recent introduction) provides a set of mathematically grounded tools for understanding evolution dynamics, at the cost of many simplifying assumptions.

Evolutionary robotics has recently gained relevance as an individual-based modeling and simulation method in evolutionary biology (Floreano and Keller, 2010; Waibel et al., 2011; Long, 2012; Mitri et al., 2013; Ferrante et al., 2015; Bernard et al., 2016), enabling the study of evolution in populations of robotic individuals in the physical world. Embodied evolution enables more accurate models of evolution because it is possible to embody not only the physical interactions but also the evolutionary operators themselves.

Synthetic approach: Embodied evolution can also be used to "understand by design" (Pfeifer and Scheier, 2001). As Maynard Smith (1992), a prominent researcher in evolutionary biology, advocated in a famous (Maynard Smith, 1992)'s Science paper (originally referring to Tierra (Ray, 1993)): "so far, we have been able to study only one evolving system and we cannot wait for interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we have to look at artificial ones."

This *synthetic approach* stands somewhere between biology and engineering, using tools from the latter to understand mechanisms originally observed in nature and aiming at identifying general principles not confined to any particular (biological) substrate. Beyond improving our *understanding* of adaptive mechanisms, these general principles can also be used to improve our ability to *design* complex systems.

6.2. Research Agenda

We identify a number of open issues that need to be addressed so that embodied evolution can develop into a relevant technique to enable online adaptivity of robot collectives. Some of these issues have been researched in other fields (e.g., credit assignment is a well-known and often considered topic in reinforcement learning research). Lessons can and should be learned from there, inspiring embodied evolution research into the relevance and applicability of findings in those other fields.

In particular, we identify the following challenges.

Benchmarks: The pseudocode in Section 3 provides a clarification of embodied evolution's concepts by describing the basic building blocks of the algorithm. This is only a first step toward a theoretical and practical framework for embodied evolution. Some authors have already taken steps in this direction. For instance, Prieto et al. (2015) propose an abstract algorithmic model to study both general and specific properties of embodied evolution implementations. Montanier et al. (2016) described "vanilla" versions of embodied evolution algorithms that can be used as practical benchmarks. Further exploration of abstract models for theoretical validation is needed. Also, standard benchmarks and test cases, along with systematically making the source code available, would provide a solid basis for empirical validation of individual contributions.

Evolutionary dynamics: Embodied evolution requires new tools for analyzing the evolutionary dynamics at work. Because the evolutionary operators apply *in situ*, the dynamics of the evolutionary process are not only important in the context of understanding or improving an optimization procedure, but they also have a direct bearing on how the robots behave and change their behavior when deployed.

Tools and methodologies to characterize the dynamics of evolving systems are available. The field of population genetics has produced techniques for estimating the selection pressure compared to genetic drift possibly occurring in finite-sized populations (see, for instance, Wakeley (2008) and Charlesworth and Charlesworth (2010) for a comprehensive introduction). Similarly, tools from adaptive dynamics (Geritz et al., 1998) can be used to investigate how particular solutions spread within the population. Finally, embodied evolution produces phylogenetic trees that can be studied either from a population genetics viewpoint (e.g., coalescence theory to understand the temporal structure of evolutionary adaptation) and graph theory (e.g., to characterize the particular structure of the inheritance graph). Boumaza (2017) shows an interesting first foray into using this technique to analyze embodied evolution.

Credit assignment: In all the research reviewed in this article that considers robot tasks, the fitness function is defined and implemented at the level of the individual robot: it assesses its own performance independently of the others. However, collectively solving a task often requires an assessment of performance at group level rather than individual level. This raises the issue of estimating each individual's contribution to the group's performance, which is unlikely to be completely captured by a fitness function (e.g. all individuals going toward the single larger food patch may not always be the best strategy if one aim to bring back the largest amount of food to the nest).

Closely related to our concern, Stone et al. (2010) formulated the *ad hoc teamwork* problem in multirobot systems, involving robots that each must "collaborate with previously unknown teammates on tasks to which they are all individually capable of contributing as team members." As stated by Wolpert and Tumer (2008), this implies devoting substantial attention to the problem of estimating the *local utility* of individual agents with respect to the *global welfare* of the whole group and how to make a tradeoff between individual and group performance (e.g., Hardin (1968); Arthur (1994)).

While a generally applicable method to estimate an individual's local utility in an online distributed setting has so far eluded the community, it is possible to provide an exact assessment in controlled settings. Methods from cooperative game theory, such as computing the Shapley value (Shapley, 1953), could be used in embodied evolution but are computationally expensive and require the ability to replay experiments. However, replaying experiments is possible only with simulation and/or controlled experimental settings. While these methods cannot apply when robots are deployed in the real world, they at least provide a method to *compare* the outcome of candidate solutions to estimate individuals' marginal contributions and choose which should be deployed.

Social complexity: Section 4 shows that embodied evolution so far demonstrated only a limited set of social organization concepts: simple cooperative and division of labor behaviors. To address more complex tasks, we must first gain a better understanding of the mechanisms required to achieve complex collective behaviors. This raises two questions. First, there is an *ethological* question: what are the behavioral mechanisms at work in complex collective behaviors? Some of them, such as positive and negative feedback between individuals, or indirect communication through the environment (i.e., *stigmergy*), are well known from examples found both in biology (Camazine et al., 2003) and theoretical physics (Deutsch et al., 2012). Second, there is a question about the origins and stability of behaviors: what are the key elements that make it possible to evolve collective behaviors, and what are their limits? Again, evolutionary ecology provides relevant insights, such as the interplay between the level of cooperation and relatedness between individuals (West et al., 2007). The literature on such phenomena in biological systems may provide a good basis for research into the evolution of social complexity in embodied evolution.

A first step would be to clearly define the nature of social complexity that is to be studied. For this, evolutionary game theory (Maynard Smith, 1992) has already produced a number of well-grounded and well-defined "games" that capture many problems involving interactions among individuals, including thorough analysis of the evolutionary dynamics in simplified setups. Of course, results obtained on abstract models may not be transferable within more realistic settings (as Bernard et al. (2016) showed for mutualistic cooperation), but the systematic use of a formal problem definition would greatly benefit the clarity of contributions in our domain.

Open-ended adaptation: As stated in Section 2, embodied evolution aims to provide continuous adaptation so that the robot collective can cope with changes in the objectives and/or the environment. Montanier and Bredeche (2011) showed that embodied evolution enables the population to react appropriately to changes in the regrowth rate of resources, but generally this aspect of embodied evolution has to date not been sufficiently addressed.

We reformulate the goal of continuous adaptation as providing *open-ended* adaptation, i.e., having the ability to continually keep exploring new behavioral patterns, constructing increasingly complex behaviors as required. Bedau et al. (2000), Soros and Stanley (2014), and Taylor et al. (2016) and others identified open-ended adaptation in artificial evolutionary systems as one of the big questions of artificial life. Open-ended adaptation in artificial systems, in particular in combination with learning relevant task behavior, has proved to be an elusive ambition.

A possible avenue to achieve this ambition may lie in the use of quality diversity approaches in embodied evolution. Recent research has considered *quality diversity* measures as a replacement (Lehman and Stanley, 2011) or additional (Mouret and Doncieux, 2012) objective to improve the population diversity and consequently the efficacy of evolution. To date, such research has focused on the evolution of behavior for particular tasks with task-specific metrics of behavioral diversity that must be tailored for each application. To be able to exploit quality diversity in unknown environments and for arbitrary tasks, generic measures of behavioral diversity must be developed.

Another avenue of research would be to take inspiration from the behavior of a passerine bird, the great tit (parus major), as recently analyzed by Aplin et al. (2017). It appears that great tits combine collective and individual learning with varying intensity as they age and that the motivations to pursue behaviors also vary with age. Reward-based learning occurs primarily in young birds and is often individual, while adult birds engage mostly in social learning to copy the behavior that is most common, regardless of whether it produces more or less rewards than alternative behavior. This combination of conformist and payoff-sensitive reinforcement allows individuals and populations both to acquire adaptive behavior and to track environmental change.

Combining embodied evolution, individual reinforcement learning with task-based and diversity-enhancing objectives may yield similar behavioral plasticity for collectives of robots.

Safety and robot ethics: To deploy the kind of adaptive technology that embodied evolution aims for responsibly, one must ensure that the adaptivity can be controlled: autonomous

adaptation carries the risk of adaptation developing in directions that do not meet the needs of human users or that they even may find undesirable. Even so, the adaptive process should be curtailed as little as possible to allow effective, open-ended, learning. The user cannot be expected to monitor and closely control the robot's behavior and learning process; this may in fact be impossible in exactly those scenarios where robotic autonomy is most beneficial and adaptivity most urgently required. There is growing awareness that it may be necessary to endow robots with innately ethical behavior (e.g., Moor (2006); Anderson and Anderson (2007); Vanderelst and Winfield (2018)), where the systems select actions based on a "moral arithmetic" (Bentham, 1878), often informed by casuistry, i.e., generalizing morality on the basis of example cases in which there is agreement concerning the correct response (Anderson and Anderson, 2007). Moral reasoning along these lines could conceivably be enabled in embodied evolution as well, in which case interactive evolution to develop surrogate models of user requirements may offer one possible route to allow user guidance.

Additional open issues and opportunities will no doubt arise from advances in this and other fields. A relevant recent development, for instance, is the possibility of evolvable morphofunctional machines that are able to change both their software *and* hardware features (Eiben and Smith, 2015) and replicate through 3D printing (Brodbeck et al., 2015). This would allow embodied evolution holistically to adapt the robots' morphologies as well as their controllers. This can have profound consequences for embodied evolution implementations that exploit these developments: it would, for instance, enable dynamic population sizes, allowing for more risky behavior as broken robots could be replaced or recycled.

REFERENCES

Alba, E. (2002). Parallel evolutionary algorithms can achieve super-linear performance. *Inf. Process. Lett.* 82, 7–13. doi:10.1016/S0020-0190(01)00281-2

Alba, E., and Dorronsoro, B. (2008). *Cellular Genetic Algorithms*. Berlin, Heidelberg, New York: Springer-Verlag.

Amato, C., Konidaris, G. G., Cruz, G., Maynor, C. A., How, J. P., and Kaelbling, L. P. (2015). "Planning for decentralized control of multiple robots under uncertainty," in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE), 1241–1248.

Anderson, M., and Anderson, S. L. (2007). Machine ethics: creating an ethical intelligent agent. *AI Mag.* 28, 15–26. doi:10.1609/aimag.v28i4.2065

Aplin, L. M., Sheldon, B. C., and McElreath, R. (2017). Conformity does not perpetuate suboptimal traditions in a wild population of songbirds. *Proc. Natl. Acad. Sci. U.S.A.* 114, 7830–7837. doi:10.1073/pnas.1621067114

Arthur, W. (1994). Inductive reasoning and bounded rationality. *Am. Econ. Rev.* 84, 406–411. doi:10.2307/2117868

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Bangel, S., and Haasdijk, E. (2017). "Reweighting rewards in embodied evolution to achieve a balanced distribution of labour," in *Proceedings of the 14th European Conference on Artificial Life ECAL 2017* (Cambridge, MA: MIT Press), 44–51.

Barrett, S., Rosenfeld, A., Kraus, S., and Stone, P. (2016). Making friends on the fly: cooperating with new teammates. *Artif. Intell.* 242, 1–68. doi:10.1016/j.artint.2016.10.005

Bayindir, L. (2016). A review of swarm robotics tasks. *Neurocomputing* 172, 292–321. doi:10.1016/j.neucom.2015.05.116

Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., et al. (2000). Open problems in artificial life. *Artif. Life* 6, 363–376. doi:10.1162/106454600300103683

7. CONCLUSION

This article provides an overview of embodied evolution for robot collectives, a research field that has been growing since its inception around the turn of the millennium. The main contribution of this article is threefold. First, it clarifies the definitions and overall process of embodied evolution. Second, it presents an overview of embodied evolution research conducted to date. Third, it provides directions for future researches.

This overview sheds light on the maturity of the field: while embodied evolution was mostly used as a parallel search method for designing individual behavior during its first decade of existence, a trend has emerged toward its collective aspects (i.e., cooperation, division of labor, specialization). This trend goes hand in hand with a trend toward larger, swarm-like, robot collectives.

We hope this overview will provide a stepping stone for the field, accounting for its maturity and acting as an inspiration for aspiring researchers. To this end, we highlighted possible applications and open issues that may drive the field's research agenda.

AUTHOR CONTRIBUTIONS

ACKNOWLEDGMENTS

Bellingham, J. G., and Rajan, K. (2007). Robotics in remote and hostile environments. *Science* 318, 1098–1102. doi:10.1126/science.1146230

Beni, G. (2005). From swarm intelligence to swarm robotics. *Robotics* 3342, 1–9. doi:10.1007/978-3-540-30552-1_1

Bentham, J. (1878). *Introduction to the Principles of Morals and Legislation*. Oxford, UK: Clarendon.

Bernard, A., André, J.-B., and Bredeche, N. (2016). To cooperate or not to cooperate: why behavioural mechanisms matter. *PLoS Comput. Biol.* 12:e1004886. doi:10.1371/journal.pcbi.1004886

Bernstein, D. S., Givan, R., Immerman, N., and Zilberstein, S. (2002). The complexity of decentralized control of Markov decision processes. *Math. Oper. Res.* 27, 819–840. doi:10.1287/moor.27.4.819.297

Bianco, R., and Nolfi, S. (2004). Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce. *Connect. Sci.* 16, 227–248. doi:10.1080/09540090412331314759

Blount, Z. D., Barrick, J. E., Davidson, C. J., and Lenski, R. E. (2012). Genomic analysis of a key innovation in an experimental *Escherichia coli* population. *Nature* 489, 513–518. doi:10.1038/nature11514

Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science* 314, 1118–1121. doi:10.1126/science.1133687

Bongard, J. C. (2013). Evolutionary robotics. *Commun. ACM* 56, 74–83. doi:10.1145/2492007.2493883

Boumaza, A. (2017). "Phylogeny of embodied evolutionary robotics," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion on – GECCO '17* (New York, NY: ACM Press), 1681–1682.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2012). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.* 7, 1–41. doi:10.1007/s11721-012-0075-2

Bredeche, N. (2014). "Embodied evolutionary robotics with large number of robots," in *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 272–273.

Bredeche, N., Haasdijk, E., and Eiben, A. E. (2009). On-line, on-board evolution of robot controllers. *Lect. Notes Comput. Sci.* 5975, 110–121. doi:10.1007/978-3-642-14156-0_10

Bredeche, N., and Montanier, J.-M. (2010). "Environment-driven embodied evolution in a population of autonomous agents," in *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature: Part II* (Berlin, Heidelberg: Springer-Verlag), 290–299.

Bredeche, N., and Montanier, J.-M. (2012). "Environment-driven open-ended evolution with a population of autonomous robots," in *Evolving Physical Systems Workshop*, 7–14.

Bredeche, N., Montanier, J.-M., and Carrignon, S. (2017). "Benefits of proportionate selection in embodied evolution: a case study with behavioural specialization," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1683–1684.

Brodbeck, L., Hauser, S., and Iida, F. (2015). Morphological evolution of physical robots through model-free phenotype development. *PLoS ONE* 10:e0128444. doi:10.1371/journal.pone.0128444

Camazine, S., Deneubourg, J.-L., Franks, N., Sneyd, J., Theraulaz, G., and Bonabeau, E. (2003). *Self-Organization in Biological Systems*. Princeton University Press.

Charlesworth, B., and Charlesworth, D. (2010). *Elements of Evolutionary Genetics*. Greenwood Village: Roberts and Company Publishers.

Christensen, D. J., Spröwitz, A., and Ijspeert, A. J. (2010). "Distributed online learning of central pattern generators in modular robots," in *From Animals to Animats 11: Proceedings of the 11th international conference on Simulation of Adaptive Behavior*, 402–412.

Cully, A., Clune, J., Tarapore, D., and Mouret, J. B. (2015). Robots that can adapt like animals. *Nature* 521, 503–507. doi:10.1038/nature14422

Deutsch, A., Theraulaz, G., and Vicsek, T. (2012). Collective motion in biological systems. *Interface Focus* 2, 689–692. doi:10.1098/rsfs.2012.0048

Dibangoye, J. S., Amato, C., Buffet, O., and Charpillet, F. (2015). "Exploiting separability in multiagent planning with continuous-state MDPs," in *IJCAI International Joint Conference on Artificial Intelligence*, 4254–4260.

Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. (2015). Evolutionary robotics: what, why, and where to. *Front. Robot. AI* 2:4. doi:10.3389/frobt.2015.00004

Eiben, A., and Smith, J. (2008). *Introduction to Evolutionary Computing*. Springer.

Eiben, A. E., Haasdijk, E., and Bredeche, N. (2010). "Embodied, on-line, on-board evolution for autonomous robotics," in *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, Vol. 5.2, eds P. Levi and S. Kernbach (Springer), 361–382.

Eiben, A. E., Schoenauer, M., Laredo, J. L. J., Castillo, P. A., Mora, A. M., and Merelo, J. J. (2007). "Exploring selection mechanisms for an agent-based distributed evolutionary algorithm," in *Proceedings of the 9th Annual Conference Companion on Genetic and Evolutionary Computation*, 2801–2808.

Eiben, A. E., and Smith, J. (2015). From evolutionary computation to the evolution of things. *Nature* 521, 476–482. doi:10.1038/nature14544

Fernandez Pérez, I., Boumaza, A., and Charpillet, F. (2014). "Comparison of selection methods in on-line distributed evolutionary robotics," in *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 1–16.

Fernandez Pérez, I., Boumaza, A., and Charpillet, F. (2015). "Decentralized innovation marking for neural controllers in embodied evolution," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 161–168.

Fernandez Pérez, I., Boumaza, A., and Charpillet, F. (2017). "Learning collaborative foraging in a swarm of robots using embodied evolution," in *Proceedings of the 14th European Conference on Artificial Life ECAL 2017* (Cambridge, MA: MIT Press), 162–161.

Ferrante, E., Turgut, A. E., Duéñez-Guzman, E., Dorigo, M., and Wenseleers, T. (2015). Evolution of self-organized task specialization in robot swarms. *PLoS Comput. Biol.* 11:e1004273. doi:10.1371/journal.pcbi.1004273

Ficici, S. G., Watson, R. A., and Pollack, J. B. (1999). "Embodied evolution: a response to challenges in evolutionary robotics," in *Proceedings of the Eighth European Workshop on Learning Robots*, eds J. L. Wyatt and J. Demiris 14–22.

Floreano, D., and Keller, L. (2010). Evolution of adaptive behaviour in robots by means of Darwinian selection. *PLoS Biol.* 8:e1000292. doi:10.1371/journal.pbio.1000292

García-Sánchez, P., Eiben, A. E., Haasdijk, E., Weel, B., and Merelo-Guervós, J.-J. J. (2012). "Testing diversity-enhancing migration policies for hybrid on-line evolution of robot controllers," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 7248 (LNCS), 52–62.

Gauci, M., Chen, J., Dodd, T. J., and Groß, R. (2012). "Evolving aggregation behaviors in multi-robot systems with binary sensors," in *Int. Symposium on Distributed Autonomous Robotic Systems (DARS 2012)* (Springer).

Geritz, S., Kisdi, E., Meszena, G., and Metz, J. (1998). Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evol. Ecol.* 12, 35–57. doi:10.1023/A:1006554906681

Good, B. H., McDonald, M. J., Barrick, J. E., Lenski, R. E., and Desai, M. M. (2017). The dynamics of molecular evolution over 60,000 generations. *Nature* 551, 45–50. doi:10.1038/nature24287

Haasdijk, E. (2015). "Combining conflicting environmental and task requirements in evolutionary robotics," in *IEEE 9th International Conference on Self-Adaptive and Self-Organizing Systems*, 131–137.

Haasdijk, E., and Bredeche, N. (2013). "Controlling task distribution in MONEE," in *Advances in Artificial Life, ECAL 2013* (MIT Press), 671–678.

Haasdijk, E., Bredeche, N., and Eiben, A. E. (2014a). Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. *PLoS ONE* 9:e98466. doi:10.1371/journal.pone.0098466

Haasdijk, E., Bredeche, N., Nolfi, S., and Eiben, A. E. (2014b). Evolutionary robotics. *Evol. Intell.* 7, 69–70. doi:10.1007/s12065-014-0113-7

Haasdijk, E., and Eigenhuis, F. (2016). "Increasing reward in biased natural selection decreases task performance," in *Artificial Life {XV}: Proceedings of the 15th International Conference on Artificial Life*, 314–322.

Haasdijk, E., Smit, S. K., and Eiben, A. E. (2012). Exploratory analysis of an on-line evolutionary algorithm in simulated robots. *Evol. Intell.* 5, 213–230. doi:10.1007/s12065-012-0083-6

Haasdijk, E., Weel, B., and Eiben, A. E. (2013). "Right on the MONEE," in *Proceeding of the Fifteenth Annual Conference on Genetic and Evolutionary Computation Conference*, ed. C. Blum (New York, NY: ACM), 207–214.

Hardin, G. (1968). The tragedy of the commons. *Science* 162, 1243–1248. doi:10.1126/science.162.3859.1243

Hart, E., Steyven, A., and Paechter, B. (2015). "Improving survivability in environment-driven distributed evolutionary algorithms through explicit relative fitness and fitness proportionate communication," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 169–176.

Hauert, S., Zufferey, J.-C., and Floreano, D. (2008). Evolved swarming without positioning information: an application in aerial communication relay. *Auton. Robot.* 26, 21–32. doi:10.1007/s10514-008-9104-9

Heinerman, J., Drupsteen, D., and Eiben, A. E. (2015). "Three-fold adaptivity in groups of robots: the effect of social learning," in *Proceedings of the 17th Annual Conference on Genetic and Evolutionary Computation, GECCO '15*, ed. S. Silva (ACM), 177–183.

Heinerman, J., Rango, M., and Eiben, A. E. (2016). "Evolution, individual learning, and social learning in a swarm of real robots," in *Proceedings – 2015 IEEE Symposium Series on Computational Intelligence, SSCI 2015* (IEEE), 1055–1062.

Hettiarachchi, S., and Spears, W. M. (2009). Distributed adaptive swarm for obstacle avoidance. *Int. J. Intell. Comput. Cybern.* 2, 644–671. doi:10.1108/17563780911005827

Hettiarachchi, S., Spears, W. M., Green, D., and Kerr, W. (2006). "Distributed agent evolution with dynamic adaptation to local unexpected scenarios," in *Innovative Concepts for Autonomic and Agent-Based Systems, Volume LNCS 3825*, eds M. G. Hinchey, P. Rago, J. L. Rash, C. A. Rouff, R. Sterritt, and W. Truszkowski (Greenbelt, MD: Springer), 245–256.

Huijsman, R.-J., Haasdijk, E., and Eiben, A. E. (2011). "An on-line on-board distributed algorithm for evolutionary robotics," in *Artificial Evolution, 10th International Conference Evolution Artificielle, Number 7401 in LNCS*, eds J.-K. Hao, P. Legrand, P. Collet, N. Monmarché, E. Lutton, and M. Schoenauer (Springer), 73–84.

Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: the use of simulation in evolutionary robotics. *Lect. Notes Comput. Sci.* 929, 704–720. doi:10.1007/3-540-59496-5_337

Karafotias, G., Haasdijk, E., and Eiben, A. E. (2011). "An algorithm for distributed on-line, on-board evolutionary robotics," in *GECCO '11: Proceedings of the 13th*

Annual Conference on Genetic and Evolutionary Computation (ACM Press), 171–178.

Kemeling, M., and Haasdijk, E. (2017). "Incorporating user feedback in embodied evolution," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion on – GECCO '17* (New York, NY: ACM Press), 1685–1686.

König, L., Mostaghim, S., and Schmeck, H. (2009). Decentralized evolution of robotic behavior using finite state automata. *Int. J. Intell. Comput. Cybern.* 2, 695–723. doi:10.1108/17563780911005845

König, L., and Schmeck, H. (2009). "A completely evolvable genotype-phenotype mapping for evolutionary robotics," in *SASO 2009 – 3rd IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, 175–185.

Lehman, J., and Stanley, K. O. (2011). Abandoning objectives: evolution through the search for novelty alone. *Evol. Comput.* 19, 189–223. doi:10.1162/EVCO_a_00025

Long, J. (2012). *Darwin's Devices: What Evolving Robots Can Teach Us about the History of Life and the Future of Technology*. Basic Books.

Mataric, M. J. (1994). *Interaction and Intelligent Behavior*. PhD thesis.

Maynard Smith, J. (1992). Evolutionary biology. Byte-sized evolution. *Nature* 355, 772–773. doi:10.1038/355772a0

Mitri, S., Wischmann, S., Floreano, D., and Keller, L. (2013). Using robots to understand social behaviour. *Biol. Rev. Camb. Philos. Soc.* 88, 31–39. doi:10.1111/j.1469-185X.2012.00236.x

Montanier, J.-M., and Bredeche, N. (2011). "Surviving the tragedy of commons: emergence of altruism in a population of evolving autonomous agents," in *European Conference on Artificial Life* (Paris, France), 550–557.

Montanier, J.-M., and Bredeche, N. (2013). "Evolution of altruism and spatial dispersion: an artificial evolutionary ecology approach," in *Advances in Artificial Life, ECAL 2013*, 260–267.

Montanier, J.-M., Carrignon, S., and Bredeche, N. (2016). Behavioural specialization in embodied evolutionary robotics: why so difficult? *Front. Robot. AI* 3:38. doi:10.3389/frobt.2016.00038

Moor, J. H. (2006). The nature, importance, and difficulty of machine ethics. *IEEE Intell. Syst.* 21, 18–21. doi:10.1109/MIS.2006.80

Mouret, J. B., and Doncieux, S. (2012). Encouraging behavioral diversity in evolutionary robotics: an empirical study. *Evol. Comput.* 20, 91–133. doi:10.1162/EVCO_a_00048

Mouret, J. B., and Tonelli, P. (2015). Artificial evolution of plastic neural networks: a few key concepts. *Stud. Comput. Intell.* 557, 251–261. doi:10.1007/978-3-642-55337-0_9

Nelson, A. L., and Grant, E. (2006). Using direct competition to select for competent controllers in evolutionary robotics. *Rob. Auton. Syst.* 54, 840–857. doi:10.1016/j.robot.2006.04.010

Nolfi, S., and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology*. Cambridge, MA: MIT Press.

Nordin, P., and Banzhaf, W. (1997). An on-line method to evolve behavior and to control a miniature robot in real time with genetic programming. *Adapt. Behav.* 5, 107–140. doi:10.1177/105971239700500201

Noskov, N., Haasdijk, E., Weel, B., and Eiben, A. E. (2013). "MONEE: using parental investment to combine open-ended and task-driven evolution," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7835 LNCS, 569–578.

Nouyan, S., Gross, R., Bonani, M., Mondada, F., and Dorigo, M. (2009). Teamwork in self-organized robot colonies. *IEEE Trans. Evol. Comput.* 13, 695–711. doi:10.1109/TEVC.2008.2011746

O'Dowd, P. J., Studley, M., and Winfield, A. F. T. (2014). The distributed co-evolution of an on-board simulator and controller for swarm robot behaviours. *Evol. Intell.* 7, 95–106. doi:10.1007/s12065-014-0112-8

Parker, L. E. (2008). "Multiple mobile robot systems," in *Handbook of Robotics*, Chap. 40, ed. S. B. Heidelberg (Springer), 921–941.

Perez, A. L. F., Bittencourt, G., and Roisenberg, M. (2008). "Embodied evolution with a new genetic programming variation algorithm," in *Proceedings – 4th International Conference on Autonomic and Autonomous Systems, ICAS 2008* (Los Alamitos, CA: IEEE Press), 118–123.

Pfeifer, R., and Scheier, C. (2001). *Understanding Intelligence*. MIT Press.

Prieto, A., Becerra, J., Bellas, F., and Duro, R. (2010). Open-ended evolution as a means to self-organize heterogeneous multi-robot systems in real time. *Rob. Auton. Syst.* 58, 1282–1291. doi:10.1016/j.robot.2010.08.004

Prieto, A., Bellas, F., and Duro, R. J. (2009). Adaptively coordinating heterogeneous robot teams through asynchronous situated coevolution. *Lect. Notes Comput. Sci.* 5864(PART 2), 75–82. doi:10.1007/978-3-642-10684-2_9

Prieto, A., Bellas, F., Trueba, P., and Duro, R. J. (2015). Towards the standardization of distributed embodied evolution. *Inf. Sci.* 312, 55–77. doi:10.1016/j.ins.2015.03.044

Prieto, A., Bellas, F., Trueba, P., and Duro, R. J. (2016). Real-time optimization of dynamic problems through distributed embodied evolution. *Integr. Comput. Aided Eng.* 23, 237–253. doi:10.3233/ICA-160522

Pugh, J., and Martinoli, A. (2009). Distributed scalable multi-robot learning using particle swarm optimization. *Swarm Intell.* 3, 203–222. doi:10.1007/s11721-009-0030-z

Ray, T. S. (1993). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artif. Life* 1, 179–209. doi:10.1162/artl.1993.1.1_2.179

Rubenstein, M., Cornejo, A., and Nagpal, R. (2014). Programmable self-assembly in a thousand-robot swarm. *Science* 345, 795–799. doi:10.1126/science.1254295

Schut, M. C., Haasdijk, E., and Prieto, A. (2009). "Is situated evolution an alternative for classical evolution?" in *IEEE Congress on Evolutionary Computation, CEC 2009*, 2971–2976.

Schwarzer, C., Hösler, C., and Michiels, N. (2010). "Artificial sexuality and reproduction of robot organisms," in *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, eds P. Levi and S. Kernbach (Berlin, Heidelberg, New York: Springer-Verlag), 384–403.

Schwarzer, C., Schlachter, F., and Michiels, N. K. (2011). Online evolution in dynamic environments using neural networks in autonomous robots. *Int. J. Adv. Intell. Syst.* 4, 288–298.

Shapley, L. S. (1953). A value for n-person games. *Ann. Math. Stud.* 28, 307–317. doi:10.1515/9781400881970-018

Silva, F., Correia, L., and Christensen, A. L. (2013). "Dynamics of neuronal models in online neuroevolution of robotic controllers," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 8154, 90–101.

Silva, F., Correia, L., and Christensen, A. L. (2017). Evolutionary online behaviour learning and adaptation in real robots. *R. Soc. Open Sci.* 4, 160938. doi:10.1098/rsos.160938

Silva, F., Duarte, M., Correia, L., Oliveira, S. M., and Christensen, A. L. (2016). Open issues in evolutionary robotics. *Evol. Comput.* 24, 205–236. doi:10.1162/EVCO_a_00172

Silva, F., Urbano, P., Correia, L., and Christensen, A. L. (2015). odNEAT: an algorithm for decentralised online evolution of robotic controllers. *Evol. Comput.* 23, 421–449. doi:10.1162/EVCO_a_00141

Silva, F., Urbano, P., Oliveira, S., and Christensen, A. L. (2012). odNEAT: an algorithm for distributed online, onboard evolution of robot behaviours. *Artif. Life* 13, 251–258. doi:10.7551/978-0-262-31050-5-ch034

Simões, E. D. V., and Dimond, K. R. K. (2001). Embedding a distributed evolutionary system into a population of autonomous mobile robots. *IEEE Int. Conf. Syst. Man Cybern.* 2, 1069–1074. doi:10.1109/ICSMC.2001.973061

Soros, L. B., and Stanley, K. K. O. (2014). "Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria," in *Proc. of Artificial Life Conference (ALife 14)*, 793–800.

Stanley, K. O., Bryant, B. D., and Miikkulainen, R. (2005). Real-time neuroevolution in the NERO video game. *IEEE Trans. Evol. Comput.* 9, 653–668. doi:10.1109/TEVC.2005.856210

Steyven, A., Hart, E., and Paechter, B. (2016). "Understanding environmental influence in an open-ended evolutionary algorithm," in *Parallel Problem Solving from Nature – PPSN XIV: 14th International Conference, Edinburgh, UK, September 17–21, 2016, Proceedings*, eds J. Handl, E. Hart, P. R. Lewis, M. López-Ibáñez, G. Ochoa, and B. Paechter (Cham: Springer), 921–931.

Stone, P., Kaminka, G. A., Kraus, S., and Rosenschein, J. S. (2010). "Ad hoc autonomous agent teams: collaboration without pre-coordination," in *The Twenty-Fourth Conference on Artificial Intelligence (AAAI)*.

Stone, P., Sutton, R. S., and Kuhlmann, G. (2005). Reinforcement learning for RoboCup-soccer keep away. *Adapt. Behav.* 13, 165–188. doi:10.1177/105971230501300301

Stone, P., and Veloso, M. (1998). Layered approach to learning client behaviors in the RoboCup soccer server. *Appl. Artif. Intell.* 12, 165–188. doi:10.1080/088395198117811

Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., et al. (2016). Open-ended evolution: perspectives from the OEE workshop in York. *Artif. Life* 22, 408–423. doi:10.1162/ARTL

Thrun, S., and Mitchell, T. M. (1995). Lifelong robot learning. *Rob. Auton. Syst.* 15, 25–46. doi:10.1016/0921-8890(95)00004-Y

Tonelli, P., and Mouret, J. B. (2013). On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks. *PLoS ONE* 8:e79138. doi:10.1371/journal.pone.0079138

Trianni, V., Nolfi, S., and Dorigo, M. (2008). "Evolution, self-organization and swarm robotics," in *Swarm Intelligence. Natural Computing Series*, eds C. Blum and D. Merkle (Berlin, Heidelberg: Springer).

Trueba, P. (2017). "Embodied evolution versus cooperative," in *Proceeding GECCO '17, Proceedings of the Genetic and Evolutionary Computation Conference Companion*, Berlin, Germany (New York, NY, USA: ACM).

Trueba, P., Prieto, A., Bellas, F., Caamaño, P., and Duro, R. J. (2011). "Task-driven species in evolutionary robotic teams," in *International Work-Conference on the Interplay between Natural and Artificial Computation* (Springer), 138–147.

Trueba, P., Prieto, A., Bellas, F., Caamaño, P., and Duro, R. J. (2012). "Self-organization and specialization in multiagent systems through open-ended natural evolution," in *Lecture Notes in Computer Science, Volume 7248 LNCS of Lecture Notes in Computer Science* (Berlin, Heidelberg: Springer), 93–102.

Trueba, P., Prieto, A., Bellas, F., Caamaño, P., and Duro, R. J. (2013). Specialization analysis of embodied evolution for robotic collective tasks. *Rob. Auton. Syst.* 61, 682–693. doi:10.1016/j.robot.2012.08.005

Urzelai, J., and Floreano, D. (2001). Evolution of adaptive synapses: robots with fast adaptive behavior in new environments. *Evol. Comput.* 9, 495–524. doi:10.1162/10636560152642887

Usui, Y., and Arita, T. (2003). "Situated and embodied evolution in collective evolutionary robotics," in *Proceedings of the 8th International Symposium on Artificial Life and Robotics, Number c*, 212–215.

Vanderelst, D., and Winfield, A. (2018). An architecture for ethical robots inspired by the simulation theory of cognition. *Cogn. Syst. Res.* 48, 56–66. doi:10.1016/j.cogsys.2017.04.002

Waibel, M., Floreano, D., and Keller, L. (2011). A quantitative test of Hamilton's rule for the evolution of altruism. *PLoS Biol.* 9:e1000615. doi:10.1371/journal.pbio.1000615

Wakeley, J. (2008). *Coalescent Theory, an Introduction*. Roberts & Company Publishers.

Walker, J. H., Garrett, S. M., and Wilson, M. S. (2006). The balance between initial training and lifelong adaptation in evolving robot controllers. *IEEE Trans. Syst. Man Cybern. B Cybern.* 36, 423–432. doi:10.1109/TSMCB.2005.859082

Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: distributing an evolutionary algorithm in a population of robots. *Rob. Auton. Syst.* 39, 1–18. doi:10.1016/S0921-8890(02)00170-7

Weel, B., Haasdijk, E., and Eiben, A. E. (2012a). "The emergence of multi-cellular robot organisms through on-line on-board evolution," in *Applications of Evolutionary Computation, Volume 7248 of Lecture Notes in Computer Science*, ed. C. Di Chio (Springer), 124–134.

Weel, B., Hoogendoorn, M., and Eiben, A. E. (2012b). "On-line evolution of controllers for aggregating swarm robots in changing environments," in *PPSN XII: Proceedings of the 12th International Conference on Parallel Problem Solving from Nature, LNCS 7492*, 245–254.

Werfel, J., Petersen, K., and Nagpal, R. (2014). Designing collective behavior in a termite-inspired robot construction team. *Science* 343, 754–758. doi:10.1126/science.1245842

West, S. A., Griffin, A. S., and Gardner, A. (2007). Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *J. Evol. Biol.* 20, 415–432. doi:10.1111/j.1420-9101.2006.01258.x

Wischmann, S., Stamm, K., and Florentin, W. (2007). "Embodied evolution and learning: the neglected timing of maturation," in *ECAL 2007: Advances in Artificial Life*, Vol. 4648 (Springer), 284–293.

Wiser, M. J., Ribeck, N., and Lenski, R. E. (2013). Long-term dynamics of adaptation in asexual populations. *Science* 342, 1364–1367. doi:10.1126/science.1243357

Wolpert, D. H., and Tumer, K. (2008). *An Introduction to Collective Intelligence*. Technical Report, NASA.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Tracking All Members of a Honey Bee Colony Over Their Lifetime Using Learned Models of Correspondence

Franziska Boenisch, Benjamin Rosemann, Benjamin Wild, David Dormagen, Fernando Wario and Tim Landgraf*

Dahlem Center for Machine Learning and Robotics, FB Mathematik und Informatik, Freie Universität Berlin, Berlin, Germany

Computational approaches to the analysis of collective behavior in social insects increasingly rely on motion paths as an intermediate data layer from which one can infer individual behaviors or social interactions. Honey bees are a popular model for learning and memory. Previous experience has been shown to affect and modulate future social interactions. So far, no lifetime history observations have been reported for all bees of a colony. In a previous work we introduced a recording setup customized to track up to 4,000 marked bees over several weeks. Due to detection and decoding errors of the bee markers, linking the correct correspondences through time is non-trivial. In this contribution we present an in-depth description of the underlying multi-step algorithm which produces motion paths, and also improves the marker decoding accuracy significantly. The proposed solution employs two classifiers to predict the correspondence of two consecutive detections in the first step, and two tracklets in the second. We automatically tracked ∼2,000 marked honey bees over 10 weeks with inexpensive recording hardware using markers without any error correction bits. We found that the proposed two-step tracking reduced incorrect ID decodings from initially ∼13% to around 2% post-tracking. Alongside this paper, we publish the first trajectory dataset for all bees in a colony, extracted from ∼3 million images covering 3 days. We invite researchers to join the collective scientific effort to investigate this intriguing animal system. All components of our system are open-source.

Keywords: honey bees, *Apis mellifera*, social insects, tracking, trajectory, lifetime history

1. INTRODUCTION

Social insect colonies are popular model organisms for self-organization and collective decision making. Devoid of central control, it often appears miraculous how orderly termites build their nests or ant colonies organize their labor. Honey bees are a particularly popular example—they stand out due to a rich repertoire of communication behaviors (von Frisch, 1965; Seeley, 2010) and their highly flexible division of labor (Robinson, 1992; Johnson, 2010). A honey bee colony robustly adapts to changing conditions, whether it may be a hole in the hive that needs to be repaired, intruders that need to be fended off, brood that needs to be reared, or food that needs to be found and processed. The colony behavior emerges from interactions of many thousand individuals. The complexity that results from the vast number of individuals is increased by the fact that bees are excellent learners: empirical evidence indicates that personal experience can modulate

communication behavior (Richter and Waddington, 1993; De Marco and Farina, 2001; Goyret and Farina, 2005; Grüter et al., 2006; Grüter and Farina, 2009; Grüter and Ratnieks, 2011; Balbuena et al., 2012). Especially among foragers, personal experience may be very variable. The various locations a forager visits might be dispersed over large distances (up to several kilometers around the hive) and each site might offer different qualities of food, or even pose threats. Thus, no two individuals share the same history and experiences. Evaluating how personal experience shapes the emergence of collective behavior and how individual information is communicated to and processed by the colony requires robust identification of individual bees over long time periods.

However, insects are particularly hard to distinguish by a human observer. Tracking a bee manually is therefore difficult to realize without marking these animals individually. Furthermore, following more than one individual simultaneously is almost impossible for the human eye. Thus, the video recording must be watched once per individual, which, in the case of a bee hive, might be several hundred or thousand times. Processing long time spans or the observation of many bees is therefore highly infeasible, or is limited to only a small group of animals. Most studies furthermore focused on one focal property, such as certain behaviors or the position of the animal. Over the last decades, various aspects of the social interactions in honey bee colonies have been investigated with remarkable efforts in data collection: Naug (2008) manually followed around 1,000 marked bees in a 1 h long video to analyze food exchange interactions. Baracchi and Cini (2014) manually extracted the positions of 211 bees once per minute for 10 h of video data to analyze the colony's proximity network. Biesmeijer and Seeley (2005) observed foraging related behaviors of a total of 120 marked bees over 20 days. Couvillon and coworkers manually decoded over 5,000 waggle dances from video (Couvillon et al., 2014). Research questions requiring multiple properties, many individuals, or long time frames are limited by the costs of manual labor.

In recent years, computer vision software for the automatic identification and tracking of animals has evolved into a popular tool for quantifying behavior (Krause et al., 2013; Dell et al., 2014). Although some focal behaviors might be extracted from the video feed directly (Berman et al., 2014; Wiltschko et al.,

2015; Wario et al., 2017), tracking the position of an animal often suffices to infer its behavioral state (Kabra et al., 2013; Eyjolfsdottir et al., 2016; Blut et al., 2017). Tracking bees within a colony is a particularly challenging task due to dense populations, similar target appearance, frequent occlusions, and a significant portion of the colony frequently leaving the hive. The exploration flights of foragers might take several hours, guard bees might stay outside the entire day to inspect incoming individuals. The observation of individual activity over many weeks, hence, requires robust means for unique identification.

For a system that robustly decodes the identity of a given detection, the tracking task reduces to simply connecting matching IDs. Recently, three marker-based insect tracking systems (Mersch et al., 2013; Crall et al., 2015; Gernat et al., 2018) have been proposed that use a binary code with up to 26 bits for error correction (Thompson, 1983). The decoding process can reliably detect and correct errors, or, reject a detection that can not be decoded. There are two disadvantages to this approach. First, error correction requires relatively expensive recording equipment (most systems use at least a 20 MP sensor with a high quality lens). Second, detections that could not be decoded can usually not be integrated into the trajectory, effectively reducing the detection accuracy and sample rate.

In contrast to these solutions, we have developed a system called BeesBook that uses much less expensive recording equipment (Wario et al., 2015). **Figure 1** shows our recording setup, **Figure 2** visualizes the processing steps performed after the recording. Our system localizes tags with a recall of 98% at 99% precision and decodes 86% IDs correctly without relying on error correcting codes (Wild et al., 2018). See **Figure 3** for the tag design. Linking detections only based on matching IDs would quickly accumulate errors, long-term trajectories would exhibit gaps or jumps between individuals. Following individuals robustly, thus, requires a more elaborate tracking algorithm.

The field of multiple object tracking has produced numerous solutions to various use-cases such as pedestrian and vehicle tracking (for reviews see Cox, 1993; Wu et al., 2013; Luo et al., 2014; Betke and Wu, 2016). Animals, especially insects, are harder to distinguish and solutions for tracking multiple animals over long time frames are far less numerous (see Dell et al., 2014 for a review on animal tracking). Since our target subjects may
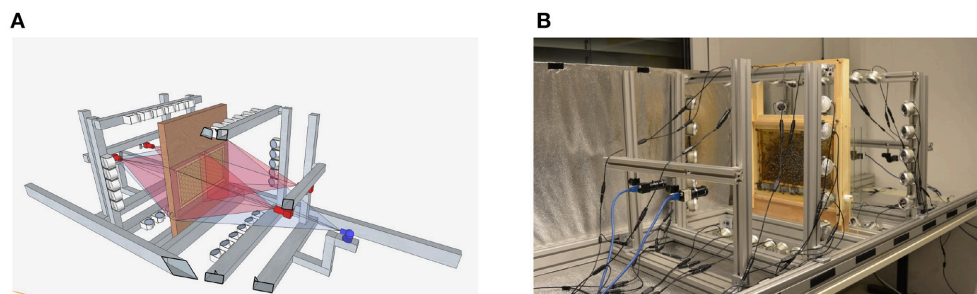
FIGURE 1 | (A) Schematic representation of the setup. Each side of the comb is recorded by two 12 MP PointGray Flea3 cameras. The pictures have an overlap of several centimeters on each side. **(B)** The recording-setup used in summer 2015. The comb, cameras, and the infrared lights are depicted, the tube that can be used by the bees to leave the setup is not visible. During recording, the setup is covered. Figures adapted from Wario et al. (2015).

FIGURE 2 | The data processing steps of the BeesBook project. The images captured by the recording setup are compressed on-the-fly to videos containing 1,024 frames each. The video data is then transferred to a large storage from where it can be accessed by the pipeline for processing. Preprocessing: histogram equalization and subsampling for the localizer. Localization: bee markers are localized using a convolutional neural network. Decoding: a second network decodes the IDs and rotation angles. Stitching: the image coordinates of the tags are transformed to hive coordinates and duplicate data in regions where images overlap are removed.

FIGURE 3 | (A) The tag-design in the BeesBook project uses 12 coding segments arranged in an arc around two semi-circles that encode the orientation of the bee. The tag is glued onto the thorax such that the white semi-circle is rotated toward the bee's head. Figure adapted from Wario (2017). **(B)** Several tagged honey bees on a comb. The round and curved tags are designed to endure heavy duty activities such as cell inspections and foraging trips.

leave the area under observation at any time, the animal's identity cannot be preserved by tracking alone. We require some means of identification for a new detection, whether it be paint marks or number tags on the animals, or identity-preserving descriptors extracted from the detection.

While color codes are infeasible with monochromatic imaging, using image statistics to fingerprint sequences of visible animals (Kühl and Burghardt, 2013; Wang and Yeung, 2013; Pérez-Escudero et al., 2014) may work even with unstructured paint markers. Merging tracklets after occlusions can then be done by matching fingerprints. However, it remains untested whether these approaches can resolve the numerous ambiguities in long-term observations of many hundreds or thousands of bees that may leave the hive for several hours.

In the following, we describe the features that we used to train machine learning classifiers to link individual detections and short tracklets in a crowded bee hive. We evaluate our results with respect to path and ID correctness. We conclude

that long-term tracking can be performed without marker-based error correction codes. Tracking can, thus, be conducted without expensive high-resolution, low-noise camera equipment. Instead, decoding errors in simple markers can be mitigated by the proposed tracking solution, leading to a higher final accuracy of the assigned IDs compared to other marker-based systems that do not employ a tracking step.

2. DESCRIPTION OF METHODS

2.1. Problem Statement and Overview of Tracking Approach

The tracking problem is defined as follows: Given a set of detections (timestamp, location, orientation, and ID information), find correct correspondences among detections over time (tracks) and assign the correct ID to each track. The ID information of the detections can contain errors. Additionally,

correct correspondences between detections of consecutive frames might not exist due to missing detections caused by occluded markers. In our dataset, the ID information consists of a number in the range of 0 to 4,095, represented by 12 bits. Each bit is given as a value between 0.0 and 1.0 which corresponds to the probability that the bit is set.

To solve the described tracking problem, we propose an iterative tracking approach, similar to previous works (for reviews, see Luo et al., 2014; Betke and Wu, 2016). We use two steps: 1. Consecutive detections are combined into short but reliable tracklets (Rosemann, 2017). 2. These tracklets are connected over longer gaps (Boenisch, 2017). Previous work employing machine learning mostly scored different distance measures separately to combine them into one thresholded value for the first tracking step (Wu and Nevatia, 2007; Huang et al., 2008; Fasciano et al., 2013; Wang et al., 2014). For merging longer tracks, boosting models to predict a ranking between candidate tracklets have been proposed (Huang et al., 2008; Fasciano et al., 2013). We use machine learning models in both steps to learn the probability that two detections, or tracklets, correspond. We train the models on a manually labeled dataset of ground truth tracklets. The features that are used to predict correspondence can differ between detection level and tracklet level, so we treat these two stages as separate learning problems. Both of our tracking steps use the Hungarian algorithm (Kuhn, 1955) to assign likely matches between detections in subsequent time steps based on the predicted probability of correspondence. In the following, we describe which features are suitable for each step and how we used various regression models to create accurate trajectories. We also explain how we integrate the ID decodings of the markers along a trajectory to predict the most likely ID for this animal, which can then be used to extract long-term tracks covering the whole lifespan of an individual. See **Figure 4** for an overview of our approach.

2.2. Step 1: Linking Consecutive Detections

The first tracking step considers detections in successive frames. To reduce the number of candidates, we consider only sufficiently close detections (we use approximately 200 pixels, or 12 mm).

From these candidate pairs we extract three features:

1. Euclidean distance between the first detection and its potential successor.
2. Angular difference of both detections' orientations on the comb plane.
3. Manhattan distance between both detections' ID probabilities.

We use our manually labeled training data to create samples with these features that include both correct and incorrect examples of correspondence. A support vector machine (SVM) with a linear kernel (Cortes and Vapnik, 1995) is then trained on these samples. We also evaluated the performance of a random forest classifier (Ho, 1995) with comparable results. We use the SVM implemented in the scikit-learn library (Pedregosa et al., 2011). Their implementation of the probability estimate uses Platt's method (Platt, 1999). This SVM can then be used get the probability of correspondence for pairs of detections that were not included in the training data. To create short tracks
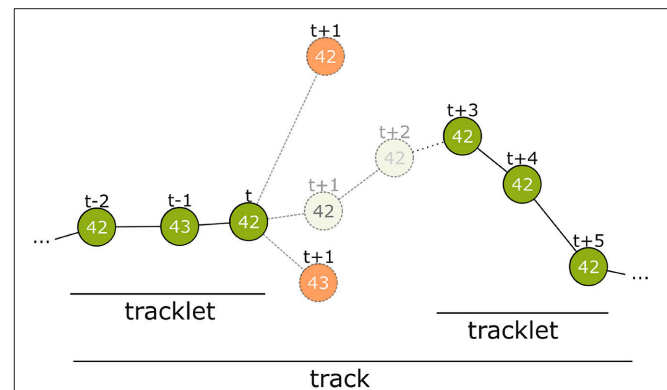

FIGURE 4 | Overview of the tracking process. The first step connects detections from successive frames to tracklets without gaps. At time step t only detections within a certain distance are considered. Even if a candidate has the same ID (top-most candidate with ID 42) it can be disregarded. The correct candidate may be detected with an erroneous ID (see t−1) or may even not be detected at all by the computer vision process. There may be close incorrect candidates that have to be rejected (candidate with ID 43 at t+1). The model assigns a correspondence probability to all the candidates. If none of them receive a sufficient score the tracklet is closed. In time step t+3 a new detection with ID 42 occurs again and is extended into a second tracklet. In tracking step 2, these tracklets are combined to a larger tracklet or track.

(tracklets), we iterate through the recorded data frame by frame and keep a list of open tracklets. Initially, we have one open tracklet for each detection of the first frame. For every time step, we use the SVM to score all new candidates against the last detection of each open tracklet. The Hungarian algorithm is then used to assign the candidate detections to the open tracklets. Tracklets are closed and not further expanded if their best candidate has a probability lower than 0.5. Detections that could not be assigned to an existing open tracklet are used to begin a new open tracklet that can be expanded in the next time step.

2.3. Step 2: Merging Tracklets

The first step yields a set of short tracklets that do not contain gaps and that could be connected with a high confidence. The second tracking step merges these tracklets into longer tracks that can contain gaps of variable duration (for distributions of tracklet and gap length in our data see section 3). Note that a tracklet could consist of a single detection or that its corresponding consecutive tracklet could still begin in the next time step without a gap. To reduce computational complexity we define a maximum gap length of 14 time steps (∼4 s in our recordings).

Similar to the first tracking step, we use the ground truth dataset to create training samples for a machine learning classifier. We create positive samples (i.e., fragments that should be classified as belonging together) by splitting each manually labeled track once at each time step. Negative samples are generated from each pair of tracks with different IDs which overlapped in time with a maximum gap size of 14. These are also split at all possible time steps. To include both more positive samples and more short track fragments in the training data, we additionally use every correct sub-track of length 3 or less

and again split it at all possible locations. This way we generated 1,021,848 training pairs, 7.4% of which were positive samples.

In preliminary tests, we found that for the given task of finding correct correspondences between tracklets, a random forest classifier performed best among a selection of classifiers available in scikit-learn (Boenisch, 2017).

Tracklets with two or more detections allow for more complex and discriminative features compared to those used in the first step. For example, matching tracklets separated by longer gaps may require features that reflect a long-term trend (e.g., the direction of motion).

We implemented 31 different features extractable from tracklet pairs. We then used four different feature selection methods from the scikit-learn library to find the features with the highest predictive power. This evaluation was done by splitting the training data further into a smaller training set and validation set. The methods used were Select-K-Best, Recursive Feature Elimination, Recursive Feature Elimination with Cross-Validation and the Random Forest Feature Importance for all possible feature subset sizes as provided by scikit-learn (Pedregosa et al., 2011). In all these methods, the same four features (number 1–4 in the listing below) performed best according to the ROC AUC score (Spackman, 1989) that proved to be a suitable metric to measure tracking results. Therefore, we chose them as an initial subset.

We then tried to improve the feature subset manually according to more tracking-specific metrics. The metrics we used were the number of tracks in the ground truth validation set that were reconstructed entirely and correctly, and the number of insertions and deletes in the tracks (for further explanation of the metrics see section 3). We added the features that lead to the highest improvements in these metrics on our validation set. This way, we first added feature 5 and then 6. After adding feature 6, the expansion of the subset with any other feature only

lead to a performance decrease in form of more insertions and less complete tracks. We therefore kept the following six features. Visualizations of features 2–5 can be found in **Figure 5**.

1. Manhattan distance of both tracklets' bitwise averaged IDs.
2. Euclidean distance of last detection of tracklet 1 to first detection of tracklet 2.
3. Forward error: Euclidean distance of linear extrapolation of last motion in first tracklet to first detection in second tracklet.
4. Backward error: Euclidean distance of linear extrapolation of first motion in second tracklet to last detection in first tracklet.
5. Angular difference of tag orientation between the last detection of the first tracklet and the first detection of the second tracklet.
6. Difference of confidence: All IDs in both tracklets are averaged with a bitwise median, we select the bit that is closest to 0.5 for each tracklet, calculate the absolute difference to 0.5 (the confidence) and compute the absolute difference of these two confidences.

2.3.1. Track ID Assignment
After the second tracking step, we determine the ID of the tracked bee by calculating the median of the bitwise ID probabilities of all detections in the track. The final ID is then determined by binarizing the resulting probabilities for each bit with probability threshold 0.5.

2.3.2. Parallelization
Tracks with a length of several minutes already display a very accurate ID decoding (see section 3). To calculate longer tracks of up to several days and weeks, we execute the tracking step 1 and step 2 for intervals of 1 h and then merge the results to longer tracks based on the assigned ID. This allows us to effectively parallelize the tracking calculation and track the entire season of
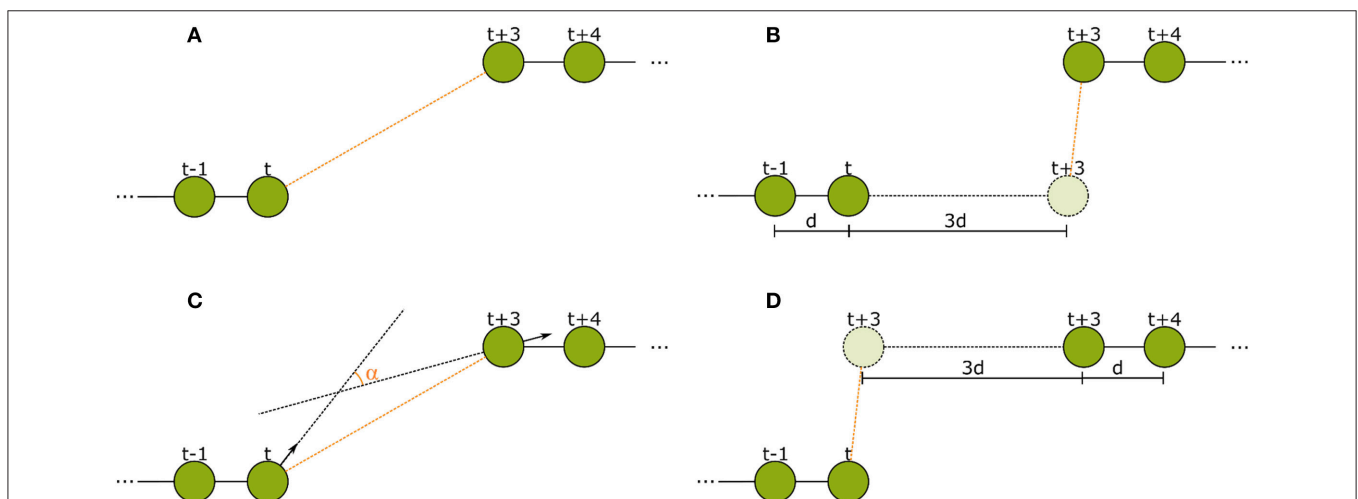
FIGURE 5 | The spatial features used in the second tracking step. **(A)** Euclidean distance between the last detection of tracklet 1 and the first detection of tracklet 2. **(B)** Forward error: Euclidean distance of the extrapolation of the last movement vector in tracklet 1 to the first detection in tracklet 2. **(C)** Angular difference between the tag orientations of the last detection in tracklet 1 and the first detection in tracklet 2. **(D)** Backward error: Euclidean distance between the reverse extrapolation of the first movement vector of tracklet 2 to the last detection of tracklet 1.

10 weeks of data in less than a week on a small cluster with <100 CPU cores.

3. RESULTS AND EVALUATION

We marked an entire colony of 1,953 bees in a 2 days session and continuously added marked young bees that were bred in an incubation chamber. In total, 2,775 bees were marked. The BeesBook system was used to record 10 weeks of continuous image data (3 Hz sample rate) of a one-frame observation hive. The image recordings were stored and processed after the recording season. The computer vision pipeline was executed on a Cray XC30 supercomputer. In total, 3,614,742,669 detections were extracted from 67,972,617 single frames, corresponding to 16,993,154 snapshots of the four cameras. Please note that the data could also be processed in real-time using consumer hardware (Wild et al., 2018).

Two ground truth datasets for the training and evaluation of our method were created manually. A custom program was used to mark the positions of an animal and to define its ID (Mischek, 2016). Details on each dataset can be found in **Table 1**. To avoid overfitting to specific colony states, the datasets were chosen to contain both high activity (around noon) and low activity (in the early morning hours) periods, different cameras and, therefore, different comb areas. Dataset *2015.1* was used to train and validate classifiers and dataset *2015.2* was used to test their performance.

Dataset *2015.1* contains 18,085 detections from which we extracted 36,045 sample pairs (i.e., all pairs with a distance of <200 pixels in consecutive frames). These samples were used to train the SVM which is used to link consecutive detections together (tracking step 1). Hyperparameters were determined manually using cross-validation on this dataset. The final model was evaluated on dataset *2015.2*.

Tracklets for the training and evaluation of a random forest classifier (tracking step 2) were extracted from datasets *2015.1*, respectively *2015.2* (see section 2 for details). Hyperparameters were optimized with hyperopt-sklearn (Komer et al., 2014) on dataset *2015.1* and the optimized model was then tested on dataset *2015.2*.

To validate the success of the tracking, we analyzed its impact on several metrics in the tracks, namely:

1. ID Improvement
2. Proportion of complete tracks
3. Correctness of resulting tracklets
4. Length of resulting tracklets

To be able to evaluate the improvement through the presented iterative tracking approach, we compare the results of the two tracking steps to the naive approach of linking the original detections over time based on their initial decoded ID only, in the following referred to as "baseline." For an overview on the improvements achieved by the different tracking steps see **Table 2**.

3.1. ID Improvement

An important goal of the tracking is to correct IDs of detections which could not be decoded correctly by the computer vision system. Without the tracking algorithm described above, all further behavioral analyses would have to consider this substantial proportion of erroneous decodings. In our dataset, 13.3% of all detections have an incorrectly decoded ID (Wild et al., 2018).

In the ground truth dataset we manually assigned detections that correspond to the same animal to one trajectory. The ground truth data can therefore be considered as the "perfect tracking." Even on these perfect tracks the median ID assignment algorithm described above provides incorrect IDs for 0.6% of all detections, due to partial occlusions, motion blur and image noise. This represents the lower error bound for the tracking system. As shown in **Figure 6**, the first tracking step reduces the fraction of incorrect IDs from 13.3 to 3.9% of all detections. The second step further improves this result to only 1.9% incorrect IDs.

TABLE 1 | Dataset *2015.1* was used for training and dataset *2015.2* for testing.

Dataset	2015.1	2015.2
Date	18.09.2015	22.09.2015
Times	11:36; 04:51	13:36
Frames	201 (3 fps)	200 (3 fps)
Detections	18,085	10,945
False positives	222 (1.23%)	82 (0.75%)
Individuals	144	98

The number of detections is the number of tags localized and decoded by the deep learning approach over all frames in the dataset. The number of false positives shows how many times the deep learning pipeline detects a detection when there is none. The number of individuals indicates how many different bees are present in the dataset.

TABLE 2 | Different metrics were used to compare the two tracking steps to both a naive baseline based on the detection IDs and to manually created tracks without errors (perfect tracking).

	Baseline (%)	After step 1 (%)	After step 2 (%)	Perfect tracking (%)
Incorrect detection IDs	13.3	3.9	1.9	0.6
Incorrect track IDs	63.5	27.2	18.2	8.2
Complete tracks	10.2	26.5	70.4	77.6
Detections missing from their track (deletions)	32.2	1.38	2.37	0
Tracks with at least one deletion	94.6	26.7	18.25	0

In all cases, the baseline performs worst and the two tracking steps successively improve the performance.

Most errors occur in short tracklets (see **Figure 7**). Therefore, the 1.9% erroneous ID assignments correspond to 18.2% of the resulting tracklets being assigned an incorrect median ID. This is an improvement over the naive baseline and the first tracking step with 63.5 and 27.2%, respectively. A perfect tracking could reduce this to 8.2% (see **Figure 8**).

3.2. Proportion of Complete Tracks
Almost all gaps between detections in our ground truth tracks are no longer than 14 frames (99.76%, see **Figure 9**). Even though
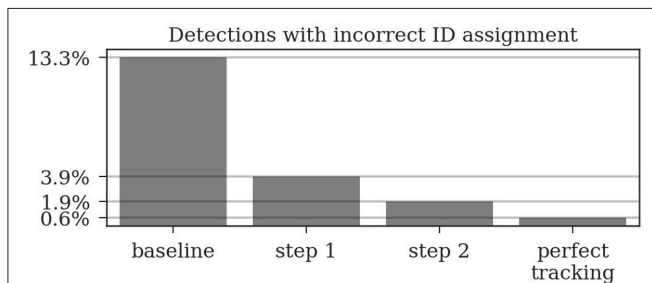

FIGURE 6 | Around 13% of the raw detections are incorrectly decoded. The first tracking step already reduces this error to around 4% and the second step further reduces it to around 2%. Even a perfect tracking (defined by the human ground truth) would still result in 0.6% incorrect IDs when using the proposed ID assignment method.
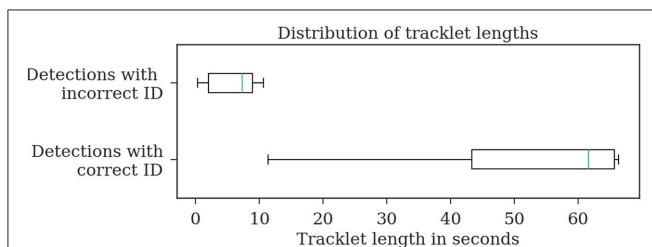

FIGURE 7 | Evaluation of the tracklet lengths of incorrectly assigned detection IDs after the second tracking step reveals that all errors in the test dataset *2015.2* happen in very short tracklets. Note that this dataset covers a duration of around 1 min.
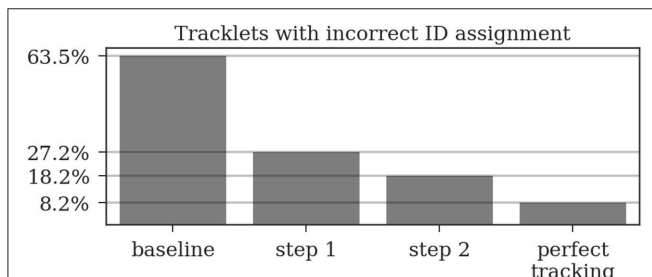

FIGURE 8 | A naive tracking approach using only the detection IDs would result in around 64% of all tracks being assigned an incorrect ID. Our two-step tracking approach reduces this to around 27 and 18%, respectively. Due to the short length of most incorrect tracklets, these 18.2% account for only 1.9% of the detections. Using our ID assignment method without any tracking errors would reduce the error to 8.2%.

large gaps between detections are rare, long tracks are likely to contain at least one such gap: Only around one third (34.7%) of the ground truth tracks contain no gaps and 77.6% contain only gaps shorter than 14 frames. As displayed in **Figure 10**, the baseline tracking finds only 10.2% complete tracks without errors (i.e., 30% of all tracks with no gaps). Step 1 is able to correctly assemble 26.5% complete tracks (i.e., around 76.5% of all tracks containing no gaps). Step 2 correctly assembles 70.4% complete tracks (about 90.4% of all tracks with a maximum gap size of <14 frames).

3.3. Correctness of Resulting Tracklets
To characterize the type of errors in our tracking results, we define a number of additional metrics. We counted detections that were incorrectly introduced into a track as *insertions*. Both tracking steps and the baseline inserted only one incorrect detection into another tracklet. Thus, <1% of both detections and tracklets were affected.

We counted detections that were missing from a tracklet (and were replaced by a gap) as *deletions*. In the baseline, 32.2% of all detections were missing from their corresponding track (94.6%
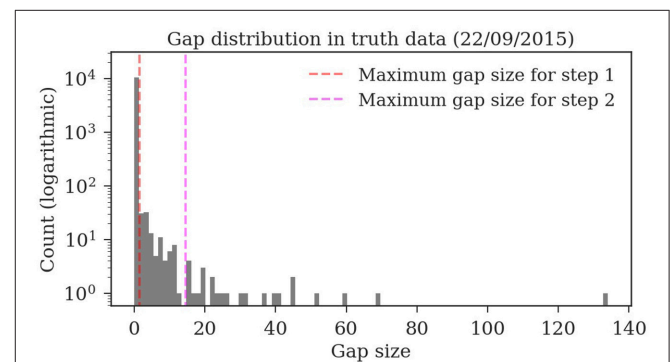
FIGURE 9 | Distribution of the gap sizes in the ground truth dataset *2015.2*. Most corresponding detections (i.e., 97.9%) have no gaps and can be therefore be matched by the first tracking step. The resulting tracklets are then merged in the second step. The maximum gap size of 14 covers 99.76% of the gaps.
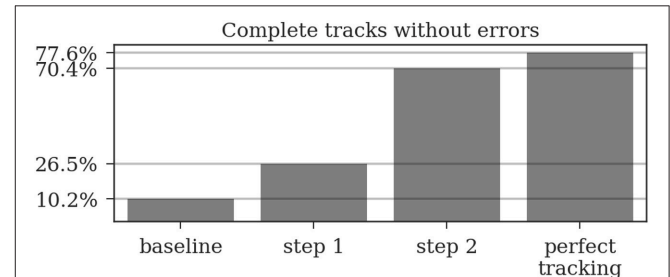
FIGURE 10 | A *complete track* perfectly reconstructs a track in our ground truth data without any missing or incorrect detections. Even a perfect tracking that is limited to a maximum gap size of 14 frames could only reconstruct around 78% of these tracks. The naive baseline based only on the detection IDs would assemble 10% without errors while our two-step tracking steps achieve 26.5 and 70.4%, respectively.

FIGURE 11 | Track lengths after tracking 1 h of video data at three frames per second. The expected length of a track is 2:23 min after the first step and 6:48 min after the second step.

of all tracks had at least one deletion). After the first step, 1.38% of detections were missing from their track, affecting 26.7% of all tracks. After the second step, 2.37% of all detections and 18.25% of all tracks were still affected.

We also evaluated whether incorrect detections were contained in a track in situations where the correct detection would have been available (instead of a gap) as *mismatches*, but no resulting tracks contained such mismatches.

3.4. Length of Resulting Tracklets

The ground truth datasets contain only short tracks with a maximum length of 1 min. To evaluate the average length of the tracks, we also tracked 1 h of data for which no ground truth data was available. The first tracking step yields shorter fragments with an expected length of 2:23 min, the second tracking step merges these fragments to tracklets with an expected duration of 6:48 min (refer to **Figure 11** for tracklet length distributions).
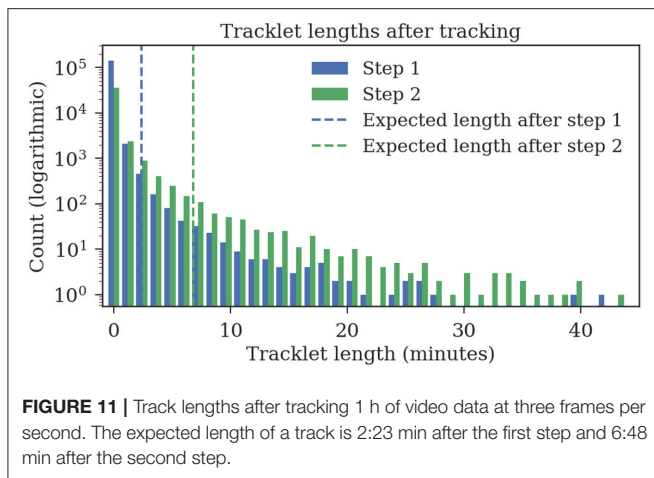
4. DISCUSSION

We have presented a multi-step tracking algorithm for fragmentary and partially erroneous detections of honey bee markers. We have applied the proposed algorithm to produce long-term trajectories of all honey bees in a colony of approximately 2,000 animals. Our dataset comprises 71 days of continuous positional data at a recording rate of 3 Hz. The presented dataset is by far the most detailed reflection of individual activities of the members of a honey bee colony. The dataset covers the entire lifespan of many hundreds of animals from the day they emerge from their brood cell until the day they die. Honey bees rely on a flexible but generally age-dependent division of labor. Hence, our dataset reflects all essential aspects of a self-sustaining colony, from an egg-laying queen and brood rearing young workers, to food collection, and colony defense. We have released a 3 days sample dataset for the interested reader (Boenisch et al., 2018). Our implementation of the proposed tracking algorithm is available online[1].

[1]https://github.com/BioroboticsLab/bb_tracking

The tracking framework presented in the previous sections is an essential part of the BeesBook system. It provides a computationally efficient approach to determine the correct IDs for more than 98% of the individuals in the honey bee hive without using extra bits for error correction.

Although it is possible to use error correction with 12 bit markers, this would reduce the number of coding bits and therefore the number of observable animals. While others chose to increase the number of bits on the marker, we solved the problem in the tracking stage. With the proposed system, we were able to reduce hardware costs for cameras and storage. When applied to the raw output of the image decoding step, the accuracy of other systems that use error-correction (for example Mersch et al., 2013) may even be improved further.

Our system provides highly accurate movement paths of bees. Given a long-term observation of several weeks, these paths, however, can still be considered short fragments. Since the IDs of these tracklets are very accurate, they can now be linked by matching IDs only.

Still, some aspects of the system can be improved. To train our classifiers, we need a sufficiently large, manually labeled dataset. Rice et al. (2015) proposed a method to create a similar dataset interactively, reducing the required manual work. Also, the circular coding scheme of our markers causes some bit configurations to appear similar under certain object poses. This knowledge could be integrated into our ID determination algorithm. The IDs along a trajectory might not provide an equal amount of information. Some might be recorded under fast motion and are therefore less reliable. Other detections could have been recorded from a still bee whose tag was partially occluded. Considering similar readings as less informative might improve the ID accuracy of our method. Still, with the proposed method there are only 1.9% detections incorrectly decoded, mostly in very short tracklets.

The resulting trajectories can now be used for further analyses of individual honey bee behavior or interactions in the social network. In addition to the three day dataset published alongside this paper, we plan to publish two more datasets covering more than 60 days of recordings, each. With this data we can investigate how bees acquire information in the colony and how that experience modulates future behavior and interactions. We hope that through this work we can interest researchers to join the collective effort of investigating the individual and collective intelligence of the honey bee, a model organism that bears a vast number of fascinating research questions.

ETHICS STATEMENT

German law does not require approval of an ethics committee for studies involving insects.

AUTHOR CONTRIBUTIONS

BR and TL: Conceptualization; FB, BR, BW, DD, and TL: Methodology; FB and BR: Software; TL: Resources, supervision,

REFERENCES

Balbuena, M. S., Molinas, J., and Farina, W. M. (2012). Honeybee recruitment to scented food sources: correlations between in-hive social interactions and foraging decisions. *Behav. Ecol. Sociobiol.* 66, 445–452. doi: 10.1007/S00265-011-1290-3

Baracchi, D., and Cini, A. (2014). A socio-spatial combined approach confirms a highly compartmentalised structure in honeybees. *Ethology* 120, 1167–1176. doi: 10.1111/eth.12290

Berman, G. J., Choi, D. M., Bialek, W., and Shaevitz, J. W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *J. R. Soc. Interface* 11:20140672. doi: 10.1098/rsif.2014.0672

Betke, M., and Wu, Z. (2016). *Data Association for Multi-Object Visual Tracking. Synthesis Lectures on Computer Vision*, (San Rafael, CA: Morgan & Claypool Publishers)

Biesmeijer, J. C., and Seeley, T. D. (2005). The use of waggle dance information by honey bees throughout their foraging careers. *Behav. Ecol. Sociobiol.* 59, 133–142. doi: 10.1007/s00265-005-0019-6

Blut, C., Crespi, A., Mersch, D., Keller, L., Zhao, L., Kollmann, M., et al. (2017). Automated computer-based detection of encounter behaviours in groups of honeybees. *Sci. Rep.* 7:17663. doi: 10.1038/s41598-017-17863-4

Boenisch, F. (2017). *Feature Engineering and Probabilistic Tracking on Honey Bee Trajectories.* Bachelor thesis, Freie Universität Berlin.

Boenisch, F., Rosemann, B., Wild, B., Wario, F., Dormagen, D., and Landgraf, T. (2018). *BeesBook Recording Season 2015 Sample.* Avaialble online at: synapse.org/#!Synapse:syn11737848

Cortes, C., and Vapnik, V. (1995). Support-vector networks. *Mach. Learn.* 20, 273–297.

Couvillon, M. J., Schürch, R., and Ratnieks, F. L. W. (2014). Waggle dance distances as integrative indicators of seasonal foraging challenges. *PLoS ONE* 9:e93495. doi: 10.1371/journal.pone.0093495

Cox, I. J. (1993). A review of statistical data association techniques for motion correspondence. *Int. J. Comput. Vis.* 10, 53–66.

Crall, J. D., Gravish, N., Mountcastle, A. M., and Combes, S. A. (2015). BEEtag: a low-cost, image-based tracking system for the study of animal behavior and locomotion. *PLoS ONE* 10:e0136487. doi: 10.1371/journal.pone.0136487

De Marco, R., and Farina, W. (2001). Changes in food source profitability affect the trophallactic and dance behavior of forager honeybees (*Apis mellifera* L.). *Behav. Ecol. Sociobiol.* 50, 441–449. doi: 10.1007/s002650100382

Dell, A. I., Bender, J. A., Branson, K., Couzin, I. D., de Polavieja, G. G., Noldus, L. P., et al. (2014). Automated image-based tracking and its application in ecology. *Trends Ecol. Evol.* 29, 417–428. doi: 10.1016/j.tree.2014.05.004

Eyjolfsdottir, E., Branson, K., Yue, Y., and Perona, P. (2016). Learning recurrent representations for hierarchical behavior modeling. arXiv:1611.00094 [cs].

Fasciano, T., Nguyen, H., Dornhaus, A., and Shin, M. C. (2013). "Tracking multiple ants in a colony," in *2013 IEEE Workshop on Applications of Computer Vision (WACV)* (Clearwater Beach, FL), 534–540.

Gernat, T., Rao, V. D., Middendorf, M., Dankowicz, H., Goldenfeld, N., and Robinson, G. E. (2018). Automated monitoring of behavior reveals bursty interaction patterns and rapid spreading dynamics in honeybee social networks. *Proc. Natl. Acad. Sci. U.S.A.* 115, 1433–1438. doi: 10.1073/pnas.1713568115

Goyret, J., and Farina, W. M. (2005). Non-random nectar unloading interactions between foragers and their receivers in the honeybee hive. *Naturwissenschaften* 92, 440–443. doi: 10.1007/s00114-005-0016-7

Grüter, C., Acosta, L. E., and Farina, W. M. (2006). Propagation of olfactory information within the honeybee hive. *Behav. Ecol. Sociobiol.* 60, 707–715. doi: 10.1007/s00265-006-0214-0

Grüter, C., and Farina, W. M. (2009). Past experiences affect interaction patterns among foragers and hive-mates in honeybees. *Ethology* 115, 790–797. doi: 10.1111/j.1439-0310.2009.01670.x

Grüter, C., and Ratnieks, F. L. W. (2011). Honeybee foragers increase the use of waggle dance information when private information becomes unrewarding. *Anim. Behav.* 81, 949–954. doi: 10.1016/j.anbehav.2011.01.014

Ho, T. K. (1995). "Random decision forests," in *Proceedings of the Third International Conference on Document Analysis and Recognition 1995, Vol. 1* (Montreal, QC), 278–282.

Huang, C., Wu, B., and Nevatia, R. (2008). "Robust object tracking by hierarchical association of detection responses," in *European Conference on Computer Vision* (Marseille: Springer), 788–801.

Johnson, B. R. (2010). Division of labor in honeybees: form, function, and proximate mechanisms. *Behav. Ecol. Sociobiol.* 64, 305–316. doi: 10.1007/s00265-009-0874-7

Kabra, M., Robie, A. A., Rivera-Alba, M., Branson, S., and Branson, K. (2013). JAABA: interactive machine learning for automatic annotation of animal behavior. *Nat. Methods* 10, 64–67. doi: 10.1038/nmeth.2281

Kühl, H. S., and Burghardt, T. (2013). Animal biometrics: quantifying and detecting phenotypic appearance. *Trends Ecol. Evol.* 28, 432–441. doi: 10.1016/j.tree.2013.02.013

Komer, B., Bergstra, J., and Eliasmith, C. (2014). "Hyperopt-sklearn: automatic hyperparameter configuration for scikit-learn," in *ICML Workshop on AutoML* (Beijing).

Krause, J., Krause, S., Arlinghaus, R., Psorakis, I., Roberts, S., and Rutz, C. (2013). Reality mining of animal social systems. *Trends Ecol. Evol.* 28, 541–551. doi: 10.1016/j.tree.2013.06.002

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Res. Logist. Q.* 2, 83–97.

Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., Zhao, X., et al. (2014). Multiple object tracking: a literature review. arXiv:1409.7618 [cs].

Mersch, D. P., Crespi, A., and Keller, L. (2013). Tracking individuals shows spatial fidelity is a key regulator of ant social organization. *Science* 340, 1090–1093. doi: 10.1126/science.1234316

Mischek, J. (2016). *Probabilistisches Tracking von Bienenpfaden.* Master thesis, Freie Universität Berlin.

Naug, D. (2008). Structure of the social network and its influence on transmission dynamics in a honeybee colony. *Behav. Ecol. Sociobiol.* 62, 1719–1725. doi: 10.1007/s00265-008-0600-x

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.

Pérez-Escudero, A., Vicente-Page, J., Hinz, R. C., Arganda, S., and de Polavieja, G. G. (2014). idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nat. Methods* 11, 743–748. doi: 10.1038/nmeth.2994

Platt, J. C. (1999). "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," in *Advances in Large Margin Classifiers*, eds A. J. Srnola, P. Bartlett, B. Schélkopf, and D. Schuurrnans (Cambridge, MA: MIT Press), 61–74.

Rice, L., Dornhausy, A., and Shin, M. C. (2015). "Efficient training of multiple ant tracking," in *IEEE Winter Conference on Applications of Computer Vision* (Waikoloa, HI), 117–123.

Richter, M. R., and Waddington, K. D. (1993). Past foraging experience influences honey bee dance behaviour. *Anim. Behav.* 46, 123–128.

Robinson, G. E. (1992). Regulation of division of labor in insect societies. *Annu. Rev. Entomol.* 37, 637–665.

Rosemann, B. (2017). *Ein Erweiterbares System für Experimente Mit Multi-Target Tracking Von Markierten Bienen.* Master thesis, Freie Universität Berlin.

Seeley, T. D. (2010). *Honeybee Democracy.* Princeton, NJ: Princeton University Press.

Spackman, K. A. (1989). "Signal detection theory: valuable tools for evaluating inductive learning," in *Proceedings of the Sixth International Workshop on Machine Learning* (San Francisco, CA: Morgan Kaufmann Publishers Inc.), 160–163.

Thompson, T. M. (1983). *From Error-Correcting Codes Through Sphere Packings to Simple Groups.* Cambridge, MA: Cambridge University Press.

von Frisch, K. (1965). *Tanzsprache und Orientierung der Bienen.* Berlin; Heidelberg: Springer.

Wang, L., Yung, N. H. C., and Xu, L. (2014). Multiple-human tracking by iterative data association and detection update. *IEEE Trans. Intell. Transport. Syst.* 15, 1886–1899. doi: 10.1109/TITS.2014.2303196

Wang, N., and Yeung, D.-Y. (2013). "Learning a deep compact image representation for visual tracking," in *Proceedings of the 26th International Conference on Neural Information Processing Systems–Vol. 1* (Portland, OR: Curran Associates Inc.), 809–817.

Wario, F. (2017). *A Computer Vision Based System for the Automatic Analysis of Social Networks in Honey Bee Colonies.* Ph.D. dissertation, Freie Universität Berlin.

Wario, F., Wild, B., Couvillon, M. J., Rojas, R., and Landgraf, T. (2015). Automatic methods for long-term tracking and the detection and decoding of communication dances in honeybees. *Front. Ecol. Evol.* 3:103. doi: 10.3389/fevo.2015.00103

Wario, F., Wild, B., Rojas, R., and Landgraf, T. (2017). Automatic detection and decoding of honey bee waggle dances. *arXiv preprint* arXiv:1708.06590.

Wild, B., Sixt, L., and Landgraf, T. (2018). Automatic localization and decoding of honeybee markers using deep convolutional neural networks. arXiv:1802.04557 [cs].

Wiltschko, A., Johnson, M., Iurilli, G., Peterson, R., Katon, J., Pashkovski, S., et al. (2015). Mapping sub-second structure in mouse behavior. *Neuron* 88, 1121–1135. doi: 10.1016/j.neuron.2015.11.031

Wu, B., and Nevatia, R. (2007). Detection and tracking of multiple, partially occluded humans by Bayesian combination of edgelet based part detectors. *Int. J. Comput. Vis.* 75, 247–266. doi: 10.1007/s11263-006-0027-7

Wu, Y., Lim, J., and Yang, M. H. (2013). "Online object tracking: a benchmark," in *2013 IEEE Conference on Computer Vision and Pattern Recognition* (Portland, OR) 2411–2418.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Predicting Dynamical Crime Distribution From Environmental and Social Influences

Simon Garnier[1]*, Joel M. Caplan[2] and Leslie W. Kennedy[2]

[1] Department of Biological Sciences, New Jersey Institute of Technology, Newark, NJ, United States, [2] School of Criminal Justice, Rutgers University-Newark, Newark, NJ, United States

Understanding how social and environmental factors contribute to the spatio-temporal distribution of criminal activities is a fundamental question in modern criminology. Thanks to the development of statistical techniques such as Risk Terrain Modeling (RTM), it is possible to evaluate precisely the criminogenic contribution of environmental features to a given location. However, the role of social information in shaping the distribution of criminal acts is largely understudied by the criminological research literature. In this paper we investigate the existence of spatio-temporal correlations between successive robbery events, after controlling for environmental influences as estimated by RTM. We begin by showing that a robbery event increases the likelihood of future robberies at and in the neighborhood of its location. This event-dependent influence decreases exponentially with time and as an inverse function of the distance to the original event. We then combine event-dependence and environmental influences in a simulation model to predict robbery patterns at the scale of a large city (Newark, NJ). We show that this model significantly improves upon the predictions of RTM alone and of a model taking into account event-dependence only when tested against real data that were not used to calibrate either model. We conclude that combining risk from exposure (past event) and vulnerability (environment), following from the Theory of Risky Places, when modeling crime distribution can improve crime suppression and prevention efforts by providing more accurate forecasting of the most likely locations of criminal events.

Keywords: crime forecasting, risk terrain modeling, event dependence, dynamical systems, vulnerability and exposure, robbery

INTRODUCTION

Recent advances in the spatial analysis of crime strongly affected the ways in which scholars and practitioners consider the origins and dispersion of crime. Hotspot mapping [1] and near repeat analysis [2] have allowed police to more efficiently target criminogenic places. Analyses of the physical contexts for crime was pioneered in criminology by Brantingham and Brantingham [3], who considered the underlying social and physical "fabric" or environmental backcloth as a framework for action. More recently, Caplan and Kennedy [4] proposed Risk Terrain Modeling (RTM) as a spatial analytical technique for empirical study of crime distribution. Resulting risk terrain maps show where certain crime events are statistically more likely to occur based on certain environmental vulnerabilities at micro places [4–7]. This technique considers the effects of multiple factors on creating distinct, identifiable areas that are conducive to crime, but emphasizes

the importance of environmental characteristics on the attraction of motivated offenders and the emergence, persistence, and desistance of crime [4–6]. For each place, it produces a risk score, that is a measure of the clustering of environmental risk factors, and can be used to forecast where crime will occur and (possibly) cluster over a period of time.

Spatial analysis—as it is used in the criminological research literature—often ignores the mechanism through which disconnected offenders cluster in space and time despite a seeming lack of deliberate coordination of activities. For instance, Sherman et al. [1] found that up to 50% of crime is produced at 3% of city locations. While Spelman [8] concluded that the statistical concentration of crime at places may be due to random and often temporary fluctuations in crime events, [9] noted that, even after correcting for such fluctuations, the worst locations accounted for a disproportionately high number of crime incidents. It appears as though, through independent action, offenders ultimately converge at the same places over given periods of time to commit similar types of crimes. If this is the case, why is this so, and how do offenders know where to go?

A possible answer comes from the concept of near repeat victimization that states that a criminal incident increases the likelihood that a nearby location or individual will be targeted in a subsequent incident [10]. This can result from either the same perpetrator repeating a crime in a location where it has been successful, or from new perpetrators encouraged directly (e.g., by a member of the same gang) or indirectly (e.g., by traces indicative of a successful event) by the first one. This has the potential of creating a positive feedback loop, with subsequent criminal events in what are defined as risky places—if close enough in space and time - increasing further the probability of additional events clustering in the same area, and so on Kennedy et al. [11]. This is described in the Theory of Risky Places [12], where the vulnerability that comes from being in high risk locations, combined with the exposure to offenders, leads to a greater probability of crime occuring.

While this concept is fairly recent in criminology, it is well-known in the scientific literature on collective behavior in biological systems. Similar feedback loops driven by past events and social information have been found to create clustering in unicellular organisms, insects, fish, birds, and mammals [13–15], even in uniform environmental conditions. However the final location of the cluster is highly dependent on the structure of the environment: clusters are more likely to originate at attractive places for the organisms, and the positive feedback process will promote the disproportionate concentration of individuals at some of the attracting places only (sometimes at a single one) while others will be abandoned [13, 16–18]. In addition, once this process has reached its stable state, the probability of starting a new cluster elsewhere—even at another attractive location - is low [13].

The striking parallel between the mechanisms of crime hotspot formation and those of clustering in social animals suggests that crime suppression and prevention efforts would strongly benefit from better understanding the combined effects of the social and physical environments in which

offenders operate. For this purpose, we propose here to combine tools for the spatial analysis of crime with methods for measuring and modeling social influence in animal groups, with the goal of improving methods for forecasting crime distribution. In particular we will use RTM as a tried and tested method to identify environmental predictors of criminal events; we will also use simulation methods to determine spatio-temporal correlations between successive events, after controlling for environmental effects. Finally, we will show that combining event-dependent and environmental influences provides improvement in forecasting changes in crime distribution over purely spatial methods (e.g., RTM) or methods based on modeling near repeat victimization only.

MATERIALS AND METHODS

Data

Crime Data

This study selectively focuses on street robberies, or robberies that occur at outdoor public spaces (e.g., streets, sidewalks, parking lots, lots/yards in front of commercial dwellings) between 2009 and 2012 in Newark, New Jersey (6,888 recorded events). The robbery data were acquired from the public records of the Newark Police Department (NPD). They only contain the time, location and nature of criminal offenses without identifying information on either the perpetrators or their victims, and therefore an ethics approval was not required as per institutional and national guidelines. Adopting the FBI's UCR Part I crime definitions, the NPD defines robbery as "the taking or attempting to take anything from the care, custody, or control of a person or persons by force or threat of force or violence and/or by putting the victim in fear" [19]. The robbery dataset includes each incident's longitude and latitude coordinates, as well as the date (e.g., 07.28.2010), day (e.g., Monday or Saturday), and hour (0–23. where 0 denotes 12 a.m.) of occurrence. For the analyses, Newark was modeled as a contiguous grid of equally sized cells the length of about half a city block (the mean blockface length is approximately 137.77 m). Each incident was therefore associated with the 68.88 m by 68.88 m cell containing its longitude and latitude coordinates.

Land Use Data

The independent variables (risk factors) of the risk terrain model were the operationalized spatial influences of land use features in Newark; the following 20 criminogenic features were included for testing in the RTM analysis: packaged liquor stores, take-out restaurants, gas stations, college campuses, parks, convenience stores, light rail stops, eat-in restaurants, foreclosed properties, parking garages, pawn shops, gyms and health clubs, grocery stores, recreation centers, at-risk housing, vacant properties, laundromats, bars, known drug markets, and schools. These data were acquired from the NPD Compstat unit or from InfoGroup, a lead provider of verified business data in the U.S.

All land use data coordinates were converted to cell coordinates matching the spatial coordinates of the crime data.

Risk Terrain Map

A risk terrain map represents the risk of a criminal event occurring at a location given the land use features of this location (see section Land Use Data above for a list of the land use features tested in this study) and relative to all the other locations considered in the analysis (all cells have approximately the size of half city blocks in Newark in this study). RTM is used to identify the relative influence of each land use feature on the occurrence of criminal events and these influences are then combined to calculate the overall relative risk associated with each considered location.

RTM has been described in detail elsewhere [6] and we will only describe its general functioning here. RTM is a two-step modeling process. In the first step RTM uses an elastic net penalization from the "penalized" R package [20] with cross-validation to perform both variable selection and regularization on a Poisson regression model of environmental risks. Model factors that stand up to shrinkage with nonzero coefficients in the penalized model are accepted as useful risk factors and passed to the next step for building the most parsimonious model.

In the second step RTM conducts a bidirectional stepwise regression using the "gamlss" R package [21] on the remaining risk factors resulting from the first step. Stepwise regression is a method to automatically reduce the complexity of a statistical model by identifying the predictive variables that significantly improve the fit to the data. The process consists in adding and removing predictive variables in a stepwise manner (i.e., one predictor at a time) and evaluating whether it significantly improves the fit to the data using in our case a BIC (Bayesian Information Criterion) score. The BIC score is a measure of the likelihood of the fit penalized by the number of predictors in the model. The model with the lowest BIC score is preferred as it strikes a balance between higher likelihood of the fit and lower complexity of the model. We repeated this process twice: once assuming a Poisson distribution of the model's residuals, and another time assuming a negative binomial distribution. Overall relative risk scores are then produced for each cell unit to produce the final risk terrain map covering the entire geographic extent of the Newark study area, which excluded the seaport and airport areas because their crimes fall under a different law enforcement jurisdiction than the NPD.

For the current study, the risk terrain map was produced using the RTMDx software, which was developed by Rutgers Center on Public Security [5]. This utility automates the RTM steps of operationalizing the spatial influence of risk factors, selecting/validating the risk factors with existing outcome event data, weighting the risk factors in relation to one another, and producing the final risk terrain map.

For each of the 20 potential risk factors described in section Land Use Data, at least 6 variables were built to measure spatial influences. These measured whether the raster cells in Newark were within 0.5, 1, 1.5, 2, 2.5, or 3 blocks of the features or in an area of high density of the factor's features. Although the extent of spatial influence can theoretically be operationalized at less than one-half block or beyond three blocks, these distances were set as the minimum and maximum search extents because they are believed to give a meaningful reach of a land use feature's

influence from a policing perspective [22, 23], and the half-block increments were used to account for varying extents of the land use features' spatial influences. For both the distance and the density calculations, we determined which cells of the study area fall into the areas defined by the different spatial extents by calculating the distance of the cell centroids to the land use feature of interest ([24], p. 5). Then, raster cells that fall within the threshold proximity (0.5, 1, 1.5, 2, 2.5, or 3 blocks) were represented as 1 (highest risk), whereas the cells outside this threshold proximity were represented as 0 (not highest risk). Density variables were reclassified into highest density (density $\geq$ mean + 2 standard deviations) and not highest density (density $<$ mean + 2 standard deviations) regions. Highest density regions were represented as 1, and regions that are not highest density were represented as 0. Ultimately, 186 model factors were produced that represent various distances from or densities of the 20 land use features in the risk terrain model. These values were then assembled into a table with rows representing cells and columns representing binary variables, and the count of street robbery events (the dependent variable) at each raster cell was calculated.

Spatio-Temporal Event-Dependence

The near repeat victimization hypothesis states that the occurrence of a criminal event at a location increases the likelihood of a subsequent event occurring at the same or a nearby location within a given time window. In order to measure this effect, we first calculate the spatio-temporal association between events as follows. For each robbery event in Newark in 2009 and 2010 we compute the probability that another event occurred within m cells from ($m = 0, 1, 2, \ldots, 40$) and n days after ($n = 0, 1, 2, \ldots, 40$) the original event.

The next step is to determine whether these probabilities are higher/lower than those expected under the assumption that there are no spatio-temporal dependence between events. For this, we use a permutation method to generate 1,000 surrogate data sets (of the same length as the original data set) in which the dependence between successive events is broken. First we randomly sample crime locations from the original data set with replacement. The probability of sampling a given location is proportional to the environmental risk value for this location as obtained from the RTM calculation (see previous section). We then associate a time to each surrogate location by randomly sampling existing occurrence times from the original data set with replacement. This procedure ensures that all surrogate events are independent in time, and that their spatial dependence is only driven by the structure of the environment, and not the location of previous events.

For each of the 1,000 surrogate data sets, we then calculate the spatio-temporal association between events following the same procedure as for the original data set. We then calculate the average ratio between the spatio-temporal association matrix of the original data set and that of the 1,000 surrogate data sets (a 2D Gaussian smoothing with standard deviation of 1 day and 1 cell is also applied to the resulting matrix). A ratio superior to 1 for a given combination of n and m indicates a likelihood higher

than random for an event to occur m cells from and n days after a previous event. A ratio inferior to 1 indicates the opposite.

If the near repeat victimization hypothesis is correct, we expect to see a maximum increase in likelihood at $m = 0$ and $n = 0$, with a progressive decrease as both m and n increase.

Forecasting Model

We propose to integrate together the environmental influences determined via RTM and the event-dependent spatio-temporal associations determined via permutation in a computer simulation model. The goal of the model is to forecast the most likely locations of future crime occurrences, given a risk terrain map locally weighted by the presence of past crime occurrences. The general functioning of the model is as follows:

1. Calculate a risk terrain map of the area of interest as described in section Risk Terrain Map. This map will be fixed for the entirety of the simulation as we assume that environmental risks do not change faster than the time scale of the simulation (a few days to a few weeks in the case of this study).
2. Calculate the spatio-temporal influence of crime occurrences on future events as described in section Spatio-Temporal Event-Dependence.
3. Create an initial spatio-temporal event-dependent risk map using existing crime occurrence data. This map should have the same dimensions as the risk terrain map. All cells of the map are initially set to 1. Each cell's value is then modified according to the distance to and time since each past crime occurrence, based on the spatio-temporal influence ratio calculated in the previous step. We assume that the respective influences of multiple past events are additive.
4. For each simulation step (1 day in the case of this study),

   a. determine the number of crime occurrences based on a distribution calculated from the original data set.
   b. For each simulated crime occurrence:

      i. Determine its location by randomly selecting a cell on the map. The probability of a cell being selected is proportional to the environmental risk value at its location on the risk terrain map multiplied by the event-dependent risk value at its location on the map.
      ii. Update the event-dependent spatio-temporal risk map to include the influence of the new simulated event.

   c. Before starting the next step, update the event-dependent spatio-temporal risk map to account for the temporal change in event-dependent spatial influence.

By simulating the model N times, we can compute a predicted probability of crime occurrence for each cell of the map.

Model Performance

We compare the forecasting performance of our model (referred to as full model in the rest of the text) against three control simulation models:

1. A random model, in which the locations of the simulated events will be selected independently of any environmental or event-dependent influence.

2. An environmental model only (referred to as RTM-only model in the rest of the text), in which the probability of a cell being selected is proportional to the environmental risk score at its location on the risk terrain map only.
3. An event-dependent spatio-temporal model only (referred to as event-only model in the rest of the text), in which the probability of a cell being selected is proportional to the the event risk value at its location on the event-dependent risk map only. Note that in this case, we will recompute the shape of the spatio-temporal influence as in section Spatio-Temporal Event-Dependence, but in absence of environmental influence.

For this comparison, we use a risk terrain map computed as described in section Risk Terrain Map using the Newark data from 2009 and 2010. The shape of the spatio-temporal influence is also computed using the 2009–2010 data. The data from 2011 and 2012 are used to initialize the event risk map and measure the performance of the models. This ensures that the model is never tested against data that has been used to parameterized it. In particular, each simulation starts at a given date in 2012 and the corresponding event risk map is initialized with all the data earlier than this date up to one year in the past.

Given a starting date, each model is simulated N times for n days after the starting date. For each actual crime occurrence in the n days after the start date, we compute the proportion ρ of simulated events that fall within 5 blocks of it. A higher average value of ρ indicates a higher clustering of simulated events around real events and therefore a better ability of the model to forecast changes in crime distribution. We can then rank the models by measuring the ratio between their average ρ and the average ρ of the random model which does not have any predictive ability.

RESULTS

Risk Terrain Map

According to the results of the bidirectional stepwise regression presented in **Table 1**, the risk terrain exhibits 11 land use features that have a criminogenic spatial influence on robberies in Newark (see **Table 1**). The Relative Risk Values in the table correspond to the exponentiated coefficients for each predictive variable in the best model selected by the RTM procedure. Once exponentiated, each coefficient is the multiplier value corresponding to a unit change in the respective predictive variable. They convey the weighting of the variables in relation to one another and reveals that a single feature might be a more or less important factor for the emergence of robberies at particular places. For instance, places influenced by nearby gas stations are almost twice as risky, or vulnerable, to robbery as places influenced by nearby takeout restaurants.

Places that are under the combined criminogenic spatial influence of these land use features had a higher risk of robberies than the places that were not. The risk terrain map represents weighted combinations of these risks at places throughout Newark, with risk scores ranging from the minimum standardized risk score of 1 to the maximum of 249.059 (see

TABLE 1 | Negative binomial type II Risk Terrain Model Factors: MuC, mean parameter coefficients; standard errors; RRV, relative risk values; Op, operationalizations; SI, spatial influence.

Risk factor	Op	SI (feet)	RRV	MuC estimate*	Std. error	*t*-value
Known drug markets	Density	226	3.058	1.1176	0.5997	18.637
Gas stations	Density	226	2.569	0.9436	0.18046	5.229
Eat-in restaurants	Density	226	2.396	0.8737	0.07152	12.217
Foreclosed properties	Proximity	678	2.227	0.8007	0.07220	11.090
Recreation centers	Proximity	452	1.911	0.6477	0.14700	4.406
Grocery stores	Proximity	1,356	1.727	0.5462	0.08602	6.349
Convenience stores	Proximity	226	1.712	0.5377	0.11250	4.799
Schools	Proximity	1,356	1.631	0.4892	0.06646	7.360
Packaged liquor stores	Density	226	1.598	0.4688	0.11755	3.988
At-risk housing	Proximity	226	1.405	0.3397	0.5349	6.352
Take-out restaurants	Proximity	1,130	1.320	0.2780	0.04880	5.696
Intercept	–	–	–	−3.0971	0.09715	−31.881

All values are significant at p < 0.001.

Figure 1A). So a place with a risk score of 249 had an expected rate of robberies that is 249 times higher than a place with a risk score of 1.

Spatio-Temporal Event-Dependence

In the 2009–2010 data, the distribution of the number of daily robbery occurrences follows a Poisson distribution with an average value of $\lambda \simeq 3.97$ (see Figure S1A) and we observe no strong autocorrelation between successive days (see Figure S1B). In the model, we use this Poisson distribution without autocorrelation to randomly allocate a number of crime occurrences to each time step of the simulation.

As expected under the near repeat victimization hypothesis, we observe a maximum increase in likelihood (3.59) at $m = 0$ cells and $n = 0$ days, with a progressive decrease as both m and n increase (see **Figure 2**). We find that this likelihood landscape can be well approximated by an inverse function of m combined with a decreasing exponential function of n, of the form:

$$1 + \frac{\alpha e^{\delta n}}{\beta m + \gamma} \qquad (1)$$

with $\alpha \simeq 2.59$, $\beta \simeq 1.47$, $\gamma = 1$ and $\delta \simeq -0.32$ when environmental influences are taken into account during the permutation process, and with $\alpha \simeq 4.62$, $\beta \simeq 1.1$, $\gamma = 1$ and $\delta \simeq -0.2$ when they are not.

We use this calibrated function in the forecasting model to generate the initial event-dependence map (see **Figure 1B** for an example) and update it after each simulated event, as described in section Forecasting Model.

Model Performance

Figure 3 shows examples of forecasting landscapes produced for January 1–7, 2012, for the RTM-only model (**Figure 3A**), the event-dependence-only model (**Figure 3B**), and the full model (**Figure 3C**), against the actual robbery occurrences during that period (white dots). **Figure 3D** shows for that particular week in 2012 how each model compares to the

random model following the procedure described in section Model Performance. All models perform better than random, with the full model combining environmental and event-dependent influences performing better than the RTM-only model accounting for environmental influences only and the event-dependence-only model accounting for spatial-temporal event dependencies only, in that order.

Figure 4 summarizes the results of similar analyzes for predictions over 3, 7, 14, and 28 days for 51 different weeks in 2012 (instead of just the first week of 2012 as in the examples in **Figure 3**). In each case, the results show that the full model performs significantly better than the random model, the event-dependence-only model and the RTM-only model, in that order (as shown by the non-overlapping notches in the boxplots; Wilcoxon ranked test, all p-values < 0.001).

Finally, **Figure 5** summarizes a direct comparison between the RTM-only model (the current state of the art in predictive criminology) and the full model that accounts for both event-dependent and environmental influences. For predictions over 1–10 days, for 51 differents weeks in 2012, the full model performs significantly better than the RTM-only model (as shown by the non-overlapping notches in the boxplots; Wilcoxon ranked test, all p-values < 0.001). Note that the variance of the data decreases with the number of days over which the predictions are calculated because of the increasing number of actual robbery occurrences that can be used to compute the average ρ value for each model.

DISCUSSION

We presented in this paper a hybrid approach to model the event-dependent and environmental drivers of criminal behaviors (more specifically robberies) at the scale of a large city (Newark, NJ). This approach combines methods from collective behavior (modeling of dynamic interactions between agents or events) and criminology (risk terrain modeling) to improve forecasting of the emergence and evolution of patterns of criminal activities. The rationale behind this hybrid approach is that RTM—the
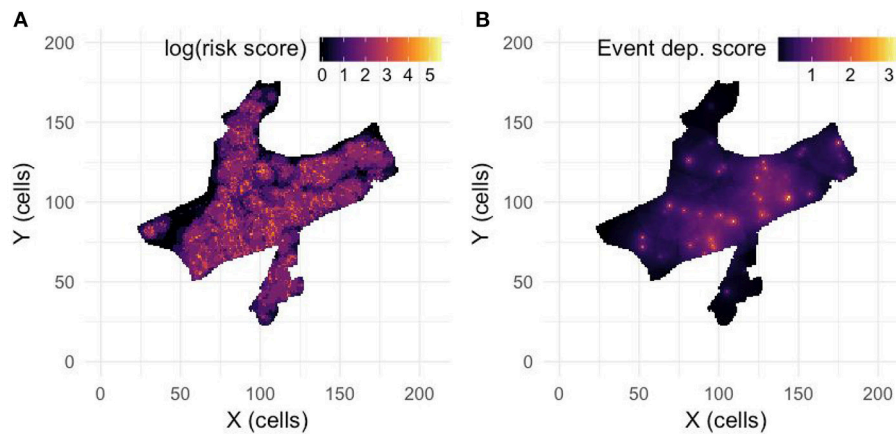
FIGURE 1 | (A) Risk terrain map calculated for robberies in Newark using data from January 1, 2009 to December 31, 2010 following the method described in section Risk Terrain Map. Lighter colors indicate cells with higher relative risk values (log scale). **(B)** Event dependence map calculated using data from January 1, 2011 to December 31, 2011 and Equation 1. Lighter colors indicate cells with elevated risk due to past robbery occurences.
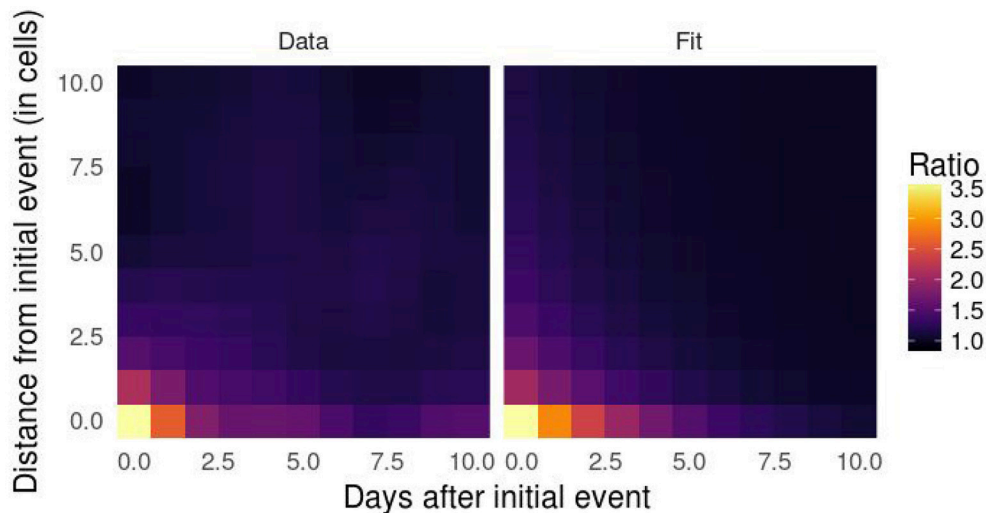
FIGURE 2 | Spatio-temporal influence of previous robbery occurrences on future events calculated with data from January 1, 2009 to December 31, 2010 following the method described in section Spatio-Temporal Event-Dependence **(Left)**, and its fit with Equation 1 **(Right)**. The colors represent the ratio between the probability that another event occurred within m cells from and n days after the original event and the same probability after random permutation of the events. As expected under the near repeat victimization hypothesis, we observe a maximum increase in likelihood (3.59) at $m = 0$ cells and $n = 0$ days, with a progressive decrease as both m and n increase.

current state-of-the-art in criminology [4]—does not account for the dependence between successive events, i.e., that the occurrence of a crime at a location increases temporarily the likelihood of future occurrences at or in the neighborhood of that location, independently of other factors such as the environmental makeup. We proposed to complement RTM with a procedure simulating this spatio-temporal event-dependence in order to obtain more accurate forecasting of the changes in the distribution of crime occurrences over time.

The first step of this procedure is to estimate the spatio-temporal dependence between successive events after controlling for environmental influences as estimated using

RTM (**Figure 1A**). Our results (**Figure 2**) show that, indeed, there is an elevated risk of robbery around previously robbed locations. This increase can be modeled as an inverse function of the distance to the original robbery, and its intensity decreases exponentially with the duration since the original robbery. This is in line with recent studies on hot spots policing that suggest that crime is not randomly distributed and is dependent on events that occurred in close proximity to new ones. For instance, based on a meta-analysis, [25] demonstrated the efficacy of event dependent approaches in increasing the chances that crime can be reduced or prevented in these areas of concentration.

FIGURE 3 | Example of predictions made by **(A)** the RTM-only model, **(B)** the event-dependent-only model, and **(C)** the full (RTM + event dependence) model, against real robbery occurrences (white dots) in Newark, from January 1, 2012 to January 7, 2012. Lighter colors indicate locations in Newark where the models predict a higher likelihood of future robbery occurrences. **(D)** Shows the improvement that each model provides over a random model without predictive ability, calculated following the procedure described in section Model Performance.

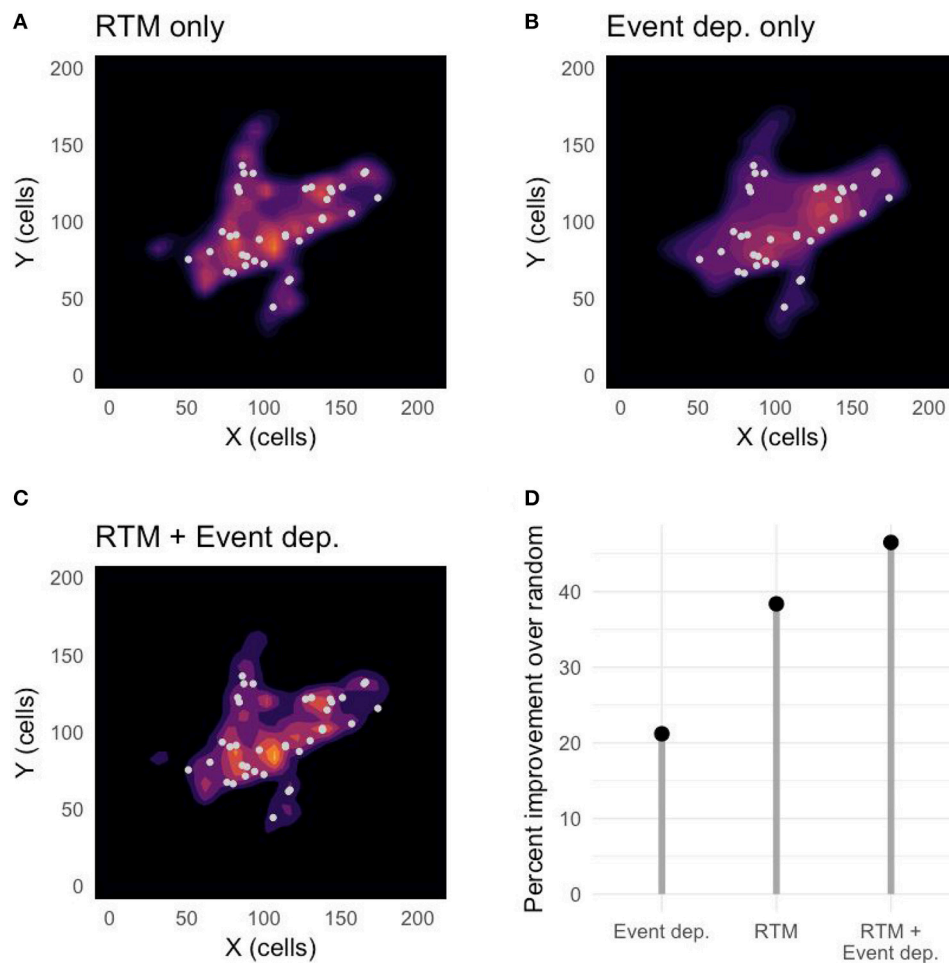The second step of the procedure uses simulations seeded with historical data to estimate the distribution of future robbery occurrences. In these simulations, the probability of a robbery happening at a location is proportional to the environmental risk at this location (as estimated by risk terrain modeling) modulated by nearby past occurrences (as estimated in the first step of the procedure). In our study, we compare the performance of this hybrid model with the performance of RTM (by setting the event-dependence to zero) and with the performance of an event-dependent only model (by setting the same environmental risk for all cells). This comparison is achieved by measuring the ability of each model to cluster predicted events for a period of time around the location of actual events during that same period of time, relative to a fully randomized model. Our results show that the predictions of the hybrid model are significantly better than the predictions of the other tested models, and that this improvement is maintained over time (at least for predictions up to 4 weeks in the future). The size of the improvement over

the RTM-only model may seem limited (4–5% in average) but it is nonetheless significant and can be explained by the quick attenuation of the spatio-temporal influence of past events (see **Figure 2**) typical of crimes of opportunity such as robberies. Larger effect sizes should be expected for crimes involving stronger interactions between the agents involved. For instance, drug markets and prostitution strolls are more enduring, often locating in the same place over long periods of time, suggesting that social factors work as facilitators in perpetuating these locations as areas of delinquency.

From a criminological perspective, our results suggest that the complexity of crime hotspots in a jurisdiction, which are derived through individual offender activities, do not necessarily require sophisticated individual behavior rules to emerge, persist, or desist. The process can be described probabilistically as a combination of environmental factors and interactions between neighboring successive events. Crimes may not always occur at expected highest-risk places or within existing hotspot areas.
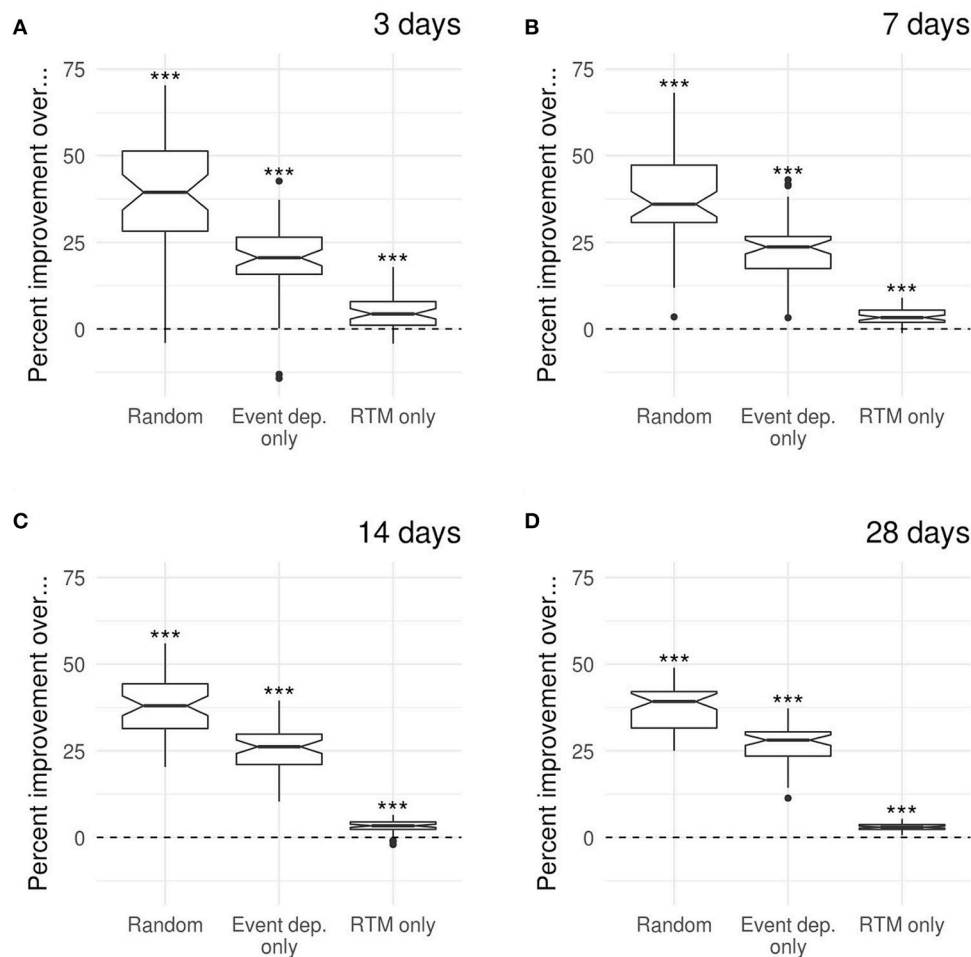
FIGURE 4 | Performance comparison between the full model and the random, event-dependent-only and RTM-only models for predictions at **(A)** 3 days, **(B)** 7 days, **(C)** 14 days, and **(D)** 28 days. Performance are shown as percent improvement of the full model over each of the other 3 models. Each boxplot corresponds to 51 measurements, each corresponding to the beginning of a different week in 2012. The notches in the boxplots correspond to the 95% confidence interval of the median and them not containing the zero line indicates strong evidence for the median to significantly differ from zero [37]. The symbols above each boxplot correspond to the significance level as calculated using a Wilcoxon ranked test with a null hypothesis of no improvement (ns, non-significant; *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$).

But, as time passes, the rational choices and stochasticity of individual offenders' decisions yields a few more crimes at the most "suitable" places. The greater number of crimes at these suitable places induces a greater number (and veracity) of perceptions that these places are "most suitable" to commit crime and reap rewards. Additional crime events stimulate more offenders to choose these places to commit their crimes, and so on Andresen [26].

So, in explaining the clustering of illegally behaving individuals, we view these as a set of dynamic mechanisms whereby hotspots appear at the global level from local interactions among its lower-level components, without being explicitly coded at the individual offender level [14]. In this scenario, positive feedback for a "hotspot cohort" of offenders results from the execution of simple behavioral "rules of thumb" that promote the creation of hotspots. A successful robbery event, for instance, whereby an offender received

cash from a victim and was never arrested or punished for it, is a kind of positive feedback which creates the conditions for similar/repeat crimes at the same locale and ultimately clustering at some places, and not others [27]. This is similar to the results of many studies on the aggregation behavior of social animals: they preferentially cluster at favorable locations but, because the individuals are also attracted toward each other, (1) they tend to aggregate at only one or a few among all the favorable locations, and (2) they can sometimes form a stable aggregate at an unfavorable location if a large enough groups has been formed there by chance [16, 18, 28]. In the criminological context, this would explain why not all high-risk locations—as predicted by RTM—become crime hotspots, and why low-risk locations may turn into hotspots in rare cases [11].

By combining environmental and event-dependent influences, our approach suggests a graduated approach to

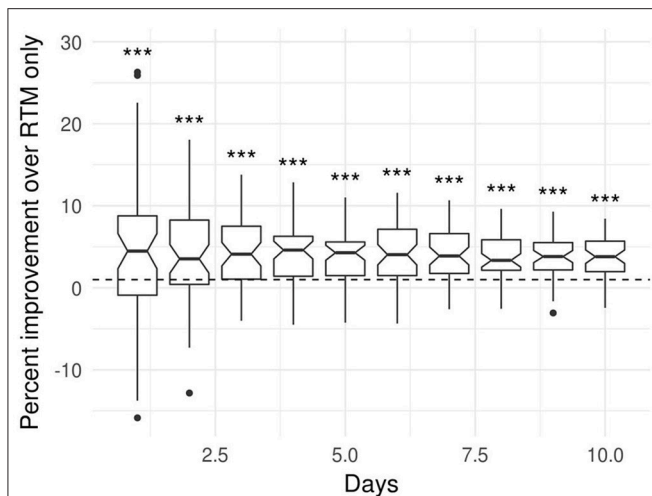**FIGURE 5 |** Performance comparison between the full model and the RTM-only model (the current state of the art model in predictive criminology) for predictions at 1–10 days. Performance are shown as percent improvement of the full model over the RTM-only model. Each boxplot corresponds to 51 measurements, each corresponding to the beginning of a different week in 2012. The notches in the boxplots correspond to the 95% confidence interval of the median and them not containing the zero line indicates strong evidence for the median to significantly differ from zero [37]. The symbols above each boxplot correspond to the significance level as calculated using a Wilcoxon ranked test with a null hypothesis of no improvement (ns, non-significant; $*p < 0.05$; $**p < 0.01$; $***p < 0.001$).

mitigating crime through intervention. At a short timescale, our model predictions can inform practitioners when allocating police resources to places forecasted to be soon in greatest need of mitigation, based on the accumulation of recent crime occurrences. This would (1) help prevent the formation of hotspots by better directing police action and (2) help identify locations to where crime might be displaced after police intervention at an emerging hotspot [29]. On a longer timescale, our ability to identify the environmental drivers of crime may help policy makers better plan the urban and economic development of neighborhoods, either avoiding environmental features that are known to increase risk, or mitigating their effect with those decreasing risk [4, 30].

CONCLUSION

Malleson et al. [31] argue that modern criminology theory has highlighted the individual-level nature of crime—whereby overall crime rates emerge from individual crimes that are committed by individual people in individual places. However, they say, "traditional modeling methodologies struggle to capture the complex dynamics of the system. The decision whether or not to commit a burglary, for example, is based on a person's unique behavioral circumstances and the immediate surrounding environment." Malleson et al. [31] add that an effective way to address these problems is through individual-level simulation

techniques such as agent-based modeling have begun to spread to the field of criminology. This paper builds on this work and provides new insights into how this approach can advance crime analysis in the future. Indeed, our work:

1. demonstrates that combining event-dependence with environmental predictors enhances our forecasts of future crime over existing methods, even in the case of crime of opportunity—such as robberies—with high attenuation rates;
2. supports the hypothesis that offenders pay attention to the results of previous crime incidences when deciding to commit a crime at a given location;
3. offers a new approach to operationalize and measure risk from exposure (past event) and vulnerability (environment) in assessing their combined spatial influence on crime outcomes and distributions;

Finally, while this approach borrows from the study of collective behaviors in biology, it reciprocates through offering a tested method to forecast behavior accounting for both individual decisions and environmental factors at different spatio-temporal scales. In addition, recent advances in understanding the role individual behavioral modulations and social networks play in shaping the collective behavior of animal groups [32–36] should provide new sources of inspiration for the design of control strategies for place-based policing and community redevelopment efforts.

AUTHOR CONTRIBUTIONS

SG, LK, and JC: conceived the study; JC: performed the RTM analysis; SG: performed the event-dependence analysis; SG: implemented all the models and performed their analysis and comparison; SG: wrote the manuscript with contributions from LK and JC.

FUNDING

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fams.2018.00013/full#supplementary-material

REFERENCES

1. Sherman LW, Gartin PR, Buerger ME. Hot spots of predatory crime: routine activities and the criminology of place. *Criminology* (1989) **27**:27–56.

2. Bowers KJ, Johnson SD. Domestic burglary repeats and space-time clusters: the dimensions of risk. *Eur J Criminol.* (2005) **2**:67–92. doi: 10.1177/1477370805048631

3. Brantingham P, Brantingham P. Criminality of place: crime generators and crime attractors. *Eur J Crim Policy Res.* (1995) **3**:5–26. Available online at: https://link.springer.com/article/10.1007%2FBF02242925

4. Caplan JM, Kennedy LW. *Risk Terrain Modeling Manual: Theoretical Framework and Technical Steps of Spatial Risk Assessment for Crime Analysis.* Newark, NJ: Rutgers Center on Public Security (2010).

5. Caplan JM, Kennedy LW, Piza EL. Joint utility of event-dependent and environmental crime analysis techniques for violent crime forecasting. *Crime Delinquency* (2012) **59**:243–70. doi: 10.1177/0011128712461901

6. Caplan JM, Kennedy LW. *Risk Terrain Modeling: Crime Prediction and Risk Reduction.* Univ of California Press (2016). Available online at: https://market.android.com/details?id=book-PinlCwAAQBAJ

7. Kennedy LW, Caplan JM, Piza E. Risk clusters, hotspots, and spatial intelligence: risk terrain modeling as an algorithm for police resource allocation strategies. *J Quant Criminol.* (2011) **27**:339–62. doi: 10.1007/s10940-010-9126-2

8. Spelman W. Abandoned buildings: magnets for crime? *J. Crim. Justice* (1993) **21**:481–95. doi: 10.1016/0047-2352(93)90033-J

9. Eck JE, Weisburd DL. Crime places in crime theory. In: John E, David E, editors. *Crime and Place.* Monsey, NY: Criminal Justice Press (1995). p. 1–33.

10. Johnson SD, Bernasco W, Bowers KJ, Elffers H, Ratcliffe J, Rengert G, et al. Space time patterns of risk: a cross national assessment of residential burglary victimization. *J Quant Criminol.* (2007) **23**:201–219. doi: 10.1007/s10940-007-9025-3

11. Kennedy LW, Caplan JM, Piza EL, Buccine-Schraeder, H. Vulnerability and exposure to crime: applying risk terrain modeling to the study of assault in chicago. *Appl Spat Anal Policy* (2015) **9**:529–48. doi: 10.1007/s12061-015-9165-z

12. Kennedy LW, Caplan JM. *A theory of Risky Places.* Newark, NJ: Rutgers Center on Public Security (2012). Available online at: http://www.rutgerscps.org/uploads/2/7/3/7/27370595/risktheorybrief_web.pdf

13. Camazine S, Deneubourg JL, Franks NR, Sneyd J, Theraulaz G, Bonabeau E. *Self-Organization in Biological Systems.* Princeton, NJ: Princeton University Press (2001).

14. Garnier S, Gautrais J, Theraulaz G. The biological principles of swarm intelligence. *Swarm Intell.* (2007) **1**:3–31. doi: 10.1007/s11721-007-0004-y

15. Sumpter DJT. The principles of collective animal behaviour. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* (2006) **361**:5–22. doi: 10.1098/rstb.2005.1733

16. Deneubourg JL, Lioni A, Detrain C. Dynamics of aggregation and emergence of cooperation. *Biol. Bull.* (2002) **202**:262–7. doi: 10.2307/1543477

17. Franks NR, Pratt SC, Mallon EB, Britton NF, Sumpter DJT. Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* (2002) **357**:1567–83. doi: 10.1098/rstb.2002.1066

18. Jeanson R, Deneubourg JL. Conspecific attraction and shelter selection in gregarious insects. *Am. Nat.* (2007) **170**:47–58. doi: 10.1086/518570

19. FBI. *Crime in the United States, (2014). U.S. Department of Justice, Federal Bureau of Investigation* (2015). Available online at: https://ucr.fbi.gov/crime-in-the-u.s/2014/crime-in-the-u.s.-2014/resource-pages/offense-definitions.pdf

20. Goeman JJ. L1 penalized estimation in the Cox proportional hazards model. *Biom. J.* (2010) **52**:70–84. doi: 10.1002/bimj.20090002

21. Stasinopoulos MD, De Bastiani F. *Flexible Regression and Smoothing: Using GAMLSS in R.* 1st edition. CRC Press/Taylor & Francis Group (2017).

22. Taylor RB. Social order and disorder of street blocks and neighborhoods: ecology, microecology, and the systemic model of social disorganization. *J. Res. Crime Delinq.* (1997) **34**:113–55.

23. Taylor RB, Harrell A. *Physical Environment and Crime.* Washington, DC: US Department of Justice, Office of Justice Programs, National Institute of Justice (1996).

24. Heffner J. Statistics of the RTMDx utility. *Risk Terrain Modeling Diagnostics Utility User Manual* (2013), p. 35–39.

25. Braga AA, Papachristos AV, Hureau DM. The effects of hot spots policing on crime: an updated systematic review and meta-analysis. *Justice Q.* (2014) **31**: 633–63. doi: 10.1080/07418825.2012.673632

26. Andresen M. A. *Environmental Criminology: Evolution, Theory, and Practice.* Taylor & Francis (2014). Available online at: https://books.google.com/books?id=8bgTAwAAQBAJ

27. Bernasco W, Block R. Robberies in Chicago: a block-level analysis of the influence of crime generators, crime attractors, and offender anchor points. *J Res Crime Delinq.* (2011) **48**:33–57. doi: 10.1177/0022427810384135

28. Canonge S, Sempo G, Jeanson R, Detrain C, Deneubourg JL. Self-amplification as a source of interindividual variability: shelter selection in cockroaches. *J. Insect Physiol.* (2009) **55**:976–82. doi: 10.1016/j.jinsphys.2009.06.011

29. Piza EL, Kennedy LW, Caplan JM. Facilitators and impediments to designing, implementing, and evaluating risk-based policing strategies using risk terrain modeling: insights from a multi-city evaluation in the United States. *Eur J Crim Policy Res.* (2018) 1–25. doi: 10.1007/s10610-017-9367-9

30. Kennedy LW, Caplan JM, Piza EL. *A Multi-Jurisdictional Test of Risk Terrain Modeling and Place-Based Evaluation of Environmental Risk-Based Patrol Deployment Strategies* (2015). Available online at: http://www.rutgerscps.org/uploads/2/7/3/7/27370595/nij6city_results_inbrief_final.pdf

31. Malleson N, Heppenstall A, See L, Evans A. Using an agent-based crime simulation to predict the effects of urban regeneration on individual household burglary risk. *Environ. Plann. B Plann. Design* (2013) **40**:405–26. doi: 10.1068/b38057

32. Becker J, Brackbill D, Centola D. Network dynamics of social influence in the wisdom of crowds. *Proc Natl Acad Sci U.S.A.* (2017) **114**:E5070–E5076. doi: 10.1073/pnas.1615978114

33. Bode NWF, Wood AJ, Franks DW. The impact of social networks on animal collective motion. *Anim. Behav.* (2011) **82**:9–38. doi: 10.1016/j.anbehav.2011.04.011

34. Michelena P, Jeanson R, Deneubourg JL, Sibbald AM. Personality and collective decision-making in foraging herbivores. *Proc R Soc B Biol Sci.* (2010) **277**:1093–1099. doi: 10.1098/rspb.2009.1926

35. Pruitt JN, Keiser CN. The personality types of key catalytic individuals shape colonies' collective behaviour and success. *Anim. Behav.* (2014) **93**:87–95. doi: 10.1016/j.anbehav.2014.04.017

36. von Rueden C, Gavrilets S, Glowacki L. Solving the puzzle of collective action through inter-individual differences. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* (2015) **370**:20150002. doi: 10.1098/rstb.2015.0002

37. Chambers JM, Cleveland WS, Kleiner B, Tukey PA. *Comparing Data Distributions.* New York, NY: Chapman & Hall(1983).

Robot Collection and Transport of Objects: A Biomimetic Process

Daniel Strömbom [1,2]* and Andrew J. King [2]

[1] Department of Mathematics, Uppsala University, Uppsala, Sweden, [2] Department of Biosciences, Swansea University, Swansea, United Kingdom

Animals as diverse as ants and humans are faced with the tasks of collecting, transporting or herding objects. Sheepdogs do this daily when they collect, herd, and maneuver flocks of sheep. Here, we adapt a shepherding algorithm inspired by sheepdogs to collect and transport objects using a robot. Our approach produces an effective robot collection process that autonomously adapts to changing environmental conditions and is robust to noise from various sources. We suggest that this biomimetic process could be implemented into suitable robots to perform collection and transport tasks that might include – for example – cleaning up objects in the environment, keeping animals away from sensitive areas or collecting and herding animals to a specific location. Furthermore, the feedback controlled interactions between the robot and objects which we study can be used to interrogate and understand the local and global interactions of real animal groups, thus offering a novel methodology of value to researchers studying collective animal behavior.

Keywords: bio-inspired robotics, feedback control, collective behavior, shepherding algorithm, adaptive system

1. INTRODUCTION

Predator attacks upon insect swarms, bird flocks, or fish schools provide a striking example of how one or a few agents (the predators) can influence the motion of many other agents (the prey) almost simultaneously (Hamilton, 1971; King et al., 2012; Handegard et al., 2012). Shepherding of sheep by dogs represents a caricature of this predator-prey interaction whereby the sheepdog maneuvers hundreds and sometimes thousands of livestock from one location to another (Strömbom et al., 2014). Engineers have long been fascinated by the act of shepherding and the behavioral rules that dogs adopt when herding since such knowledge may have application to engineering tasks as diverse as guiding groups of exploring robots (Turgut et al., 2008) to cleaning up the environment (Fingas, 2016). To this end, Strömbom et al. (2014) designed a general shepherding algorithm inspired by empirical data collected from real-life sheepdog interactions; it was proposed that the algorithm could support the efficient design of robots herding autonomous agents in a variety of contexts.

Research with multi-robot systems have sought to bring objects (and other robots) in the environment together as quickly as possible, into one cluster (Melhuish et al., 2001; Gauci et al., 2014), and such "herding" robot systems could have the potential to limit the spread of oil spills in the oceans (Zahugi et al., 2012; Fingas, 2016), and to collect rubbish (Bonnema, 2012), specific objects (Karunasena et al., 2008), or hazardous material (Nguyen et al., 2002) on both land and water. Whilst a large number of algorithms have been proposed for use in such tasks (Lien et al., 2004, 2005; Miki and Nakamura, 2006; Bennett and Trafankowski, 2012; Strömbom et al., 2014) most are studied via simulation and only capable of collecting or herding relatively low

numbers of objects or agents, at least when only one shepherd is used (Bennett and Trafankowski, 2012). The use of robots for collection and herding objects in the real-world therefore remains rare, and herding free-living animals presents an even greater challenge, given that prey animals have evolved a variety of mechanisms to avoid detection and capture (Ioannou et al., 2012). In fact, the only published research we know to successfully apply a robot for herding free-living animals is work by Vaughan et al. (1998) who designed and used a robot to herd flocks of ducks.

Introducing robots into animal groups to influence/study the behavior of the animals has been much more common (and successful) in the field of collective animal behavior (Krause et al., 2011). Robots have been used to study the behavior of cockroaches (e.g. Halloy et al., 2007), fish (e.g. Faria et al., 2010; Swain et al., 2012; Landgraf et al., 2013, 2016; Cazenille et al., 2017) and rats (e.g. Shi et al., 2013). In most cases the interactions between the animals and the robot are essentially one-way; the animals are influenced by the robot but the robot is not directly influenced by the animals. However, examples do exist where two-way interactions between a robot and a group of animals are achieved. For example, in Swain et al. (2012) a feedback controlled robot-fish interacts with a school of free-moving fish in real time. The robot fish was programmed to chase the centroid of the fish school and dart towards them when their polarization was close to zero (milling or disordered school). Such examples demonstrate the potential for using robot-animal interactions, but to fully utilize robots in the study of collective behavior, the robots need to be able to respond to the real-life individuals (and not just the collective), in real-time (Krause et al., 2011).

To advance the study and analysis of robot-animal interactions requires an integrated design process (Hamann et al., 2016) that affords remotely controlled robots and 2d or 3d tracking of robot and object/animals. The task of fully automating the tracking of multiple objects can be "surprisingly problematic under experimental conditions" (Krause et al., 2011) but advances in image tracking technologies especially via open-source software (e.g. Pérez-Escudero et al., 2014) is making this more achievable. For example, the use of a surveillance drone providing a shepherding robot with information in real time about target objects or animals would revolutionize numerous cleanup processes, and enable robots to respond to their targets even when these targets are mobile or unpredictable in some way.

Here, we present an adaptive collection robot that is part of a feedback-controlled image-based tracking system designed to target and retrieve objects. The robot algorithm is an adapted version of Strömbom et al.'s (2014) bio-inspired model of shepherding behavior that matched empirical data collected with a sheepdog and sheep in the real-world when analyzed via computer simulations. We take the Strömbom et al. (2014) algorithm, modify it, and implement it in a single robot shepherd that collects and moves objects to a given location based on feedback from the image-based tracking. We demonstrate the collection capabilities of the robot in fixed and changing environments, and show that it is fully adaptive, robust to various sources of noise, and mimicks the sheepdog behavior on which it is based. We also explain why we believe that our algorithm is a viable candidate for implementation into suitable robots to collect and move living and artificial object in the real world and, crucially, how it can also be useful to study collective animal behavior via robots.

FIGURE 1 | (A) Photo of the arena setup showing the white arena floor, red objects, black robot controlled magnet, overhead camera, and the computer used to coordinate and run all parts of the feedback control loop. Inset: e-puck robot fitted with a large red magnet used to connect to the small black magnet moving on the arena floor. **(B)** Schema illustrating the experimental setup.
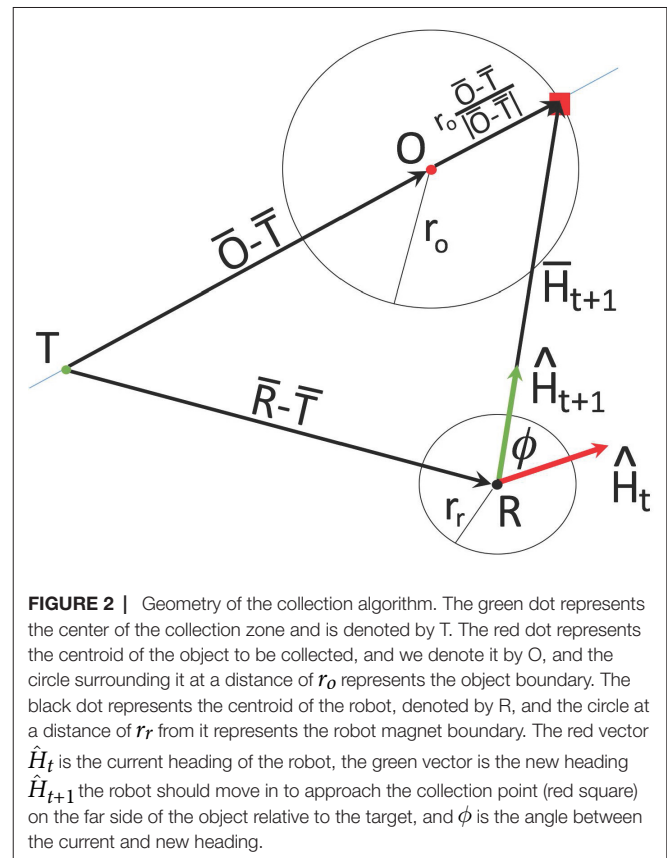
2. MATERIAL AND METHODS

2.1. Test Arena

We use an arena setup (**Figure 1A,B**) and feedback control loop similar to those employed in Swain et al. (2012) and Bonnet et al. (2017) to explore the capacity and behavior of an adapted shepherding algorithm implemented in an e-puck robot (Mondada et al., 2009) instructed to adaptively collect objects scattered in the arena to a designated collection zone. The arena floor is made out of acrylic and boundaries of the same material have been set up limiting the space available for herding to 880 × 435 mm. In our set-up, the robot moves under the arena floor and controls the movement of a black magnet (radius 5 mm) which interacts with red round objects (radius 15 mm) via physical contact on the arena floor. The robot is connected to the computer via Bluetooth and instructions are transmitted to it via the e-puck Matlab control application ePic2 (Hubert and Weibel, 2008).

2.2. Feedback Control Loop

We use an overhead camera (Logitech C902 HD pro USB) linked to a computer running Matlab R2015b. The camera takes an image, which is processed and the coordinates and radii of the objects and the robot controlled magnet are extracted using elementary image processing and analysis. The current (time) normalized orientation/heading of the robot $\hat{H}_t$, radius of the objects (r_o), and the radius of the robot magnet (r_r) are also calculated. The centroid coordinates and the radii of the objects and robot magnet are then used to calculate a new robot heading $\hat{H}_{t+1}$ for the next time step using the shepherding algorithm which is described in 2.2.2 below. The process continues until all N objects have been delivered to the collection zone which is a discshaped region near the center of the arena with radius $4 + 0.5N^{2/3}r_o$ (**Figure 1B**).

2.2.1. Image Processing and Analysis

We chose to use red objects, a black robot controlled magnet, and a white arena floor because this enabled fast, low-level image processing analysis methods on low resolution images (640 × 480). Here we describe the steps involved in the image analysis, and when applicable, include the Matlab command used in parenthesis following the description. Once an image has been imported to Matlab we overexpose it slightly and then segment the black and red objects by simple thresholding. A morphological operation is then applied to fill any "holes" in the segmented objects (*imfill*) and the centroids of the segmented objects are then calculated (*regionprops centroid*). Finally, the areas of the robot magnet and an object in the image are estimated by counting object pixels in the segmented images (*nnz*) and from these areas the radius of the robot magnet r_r and the radius r_o of the objects are calculated. As the objects do not change size the radii are only calculated on the first time step of each trial. At the beginning of each trial the current heading (in arena coordinates) $\hat{H}_0$ of the robot is estimated by extracting the centroids of the robot magnet in two successive webcam photos, acquired while the robot is moving straight ahead in its local coordinate system.

FIGURE 2 | Geometry of the collection algorithm. The green dot represents the center of the collection zone and is denoted by T. The red dot represents the centroid of the object to be collected, and we denote it by O, and the circle surrounding it at a distance of r_o represents the object boundary. The black dot represents the centroid of the robot, denoted by R, and the circle at a distance of r_r from it represents the robot magnet boundary. The red vector $\hat{H}_t$ is the current heading of the robot, the green vector is the new heading $\hat{H}_{t+1}$ the robot should move in to approach the collection point (red square) on the far side of the object relative to the target, and ϕ is the angle between the current and new heading.

2.2.2. The Shepherding Algorithm

The shepherding algorithm is modified from the collection part of the algorithm in Strömbom et al. (2014), adapting it for use with non-self-propelled objects with contact repulsion. The algorithm is designed to collect the object furthest away from the collection zone first, unless it is already in contact with another object, in which case it delivers that object to the collection zone first before venturing out towards the furthest away object. **Figure 2** illustrates how the new robot heading $\hat{H}_{t+1}$ is calculated from known quantities once a specific object has been selected for collection. We use hat notation for unit vectors and bar notation for non-normalized vectors. T denotes the center of the collection zone, O the centroid of the object to be collected, and R the centroid of the robot. The new heading of the robot is set towards the point on the object boundary on the far side of the centroid of the object O relative to the target T. This point is represented by a red square in **Figure 2** and we see that the new heading vector from the robot towards this point is given by

$$\bar{H}_{t+1} = \left(\bar{O} - \bar{T} \right) + r_o \frac{\bar{O} - \bar{T}}{|\bar{O} - \bar{T}|} - \left(\bar{R} - \bar{T} \right) = \bar{O} - \bar{R} + r_o \frac{\bar{O} - \bar{T}}{|\bar{O} - \bar{T}|}. \quad (1)$$

Once the algorithm has calculated a new heading $\bar{H}_{t+1}$ for the robot, the signed angle ϕ between the normalized current heading $\hat{H}_t$ and the normalized new heading $\hat{H}_{t+1}$ is calculated (**Figure 2**). If the magnitude of this angle is smaller than a specified threshold (0.25 radians $\approx$ 14 degrees), the robot controller instructs the robot to keep moving forward, otherwise the controller rotates the robot into

alignment with the new heading before moving forward. Once the robot has been moved the loop starts over and a new photo is taken by the overhead webcam. This process continues until all objects present on the arena floor has been delivered to a pre-assigned collection zone.

2.2.3. Experiments

We conducted a series of experiments to investigate the collection capacity and behavior of the robot. We examined situations with a fixed number of objects to be collected (phase one) and situations where the number of objects changed over time (phase two). In phase one, we ran four trials each with 2, 4, 8 and 16 objects. Objects were distributed in the arena so that no object was in the collection zone or touching an arena boundary initially. In phase two, three trials were conducted, and in each case the number of objects for collection increased within the trial. Phase two trials started with two objects, and then we added two more, then four more, and finally eight more. Objects were added to the arena once the robot was driving the final object in the arena (i.e., the 2nd, 4th, and 8th object) towards the collection zone. Trials where objects tossed into the arena ended up in the collection zone were excluded. Across all trials (both phases) the robot always started near the center of the arena and each trial terminated when all objects had been delivered to the collection zone. We collected the coordinates of the robot and the objects throughout the trials and the time to completion of each trial was recorded.

2.2.4. Measures

To evaluate the collection capacity of the robot and characterize the collection process we constructed time series with (i) the mean object-target distances, and (ii) the area occupied (convex hull) by objects. To evaluate the behavioral mechanisms by which the robot herded and collected objects we also recorded (iii) where the robot was located relative to the position of the object being herded and final target destination. To this end, we expressed the coordinates

of the robot centroid in a coordinate system that is centered on the centroid of the closest object and in which the direction towards the target is the positive x-axis. More specifically, on each time step we determine if the robot is within a distance of $2r_O$ (our definition of close) from any object and if so proceed with steps 1–3 below.

1. First specify that the centroid of the object O is the origin of the new coordinate system and then translate the robot centroid R and target coordinate T accordingly. That is, $R' = R - O$ and $T' = T - O$.
2. Calculate the (signed) angle θ of the origin to target vector.
3. Rotate the translated robot vector R' and the translated target vector T' by θ.

3. RESULTS

3.1. Robot Performance in Task

Examples of the robot collection process are provided in Video S1 which shows one collection trial each for 2, 4, 8 and 16 objects, and one trial with an increasing number of objects. All objects and zones shown in Video S1 have been superimposed on the webcam image: Target (blue asterisk), Collection zone (green ring), Object centroids (red asterisks), Object boundaries (red ring), Robot controlled magnet centroid (black asterisk), Robot controlled magnet boundary (black ring), Current heading (red rod), New (ideal) heading (green rod). The mean average distance of objects to the collection zone (**Figure 3**), and the dispersion of the objects as described by a convex hull (**Figure 3C,D**) during trials illustrate the performance of the robot for fixed and variable number of object trials.

Figure 3A We found the completion times across trials for a fixed number of objects were similar (**Figure 3A**) and mean

FIGURE 3 | (A–B) Mean object-collection zone distance over time. Thin lines show the mean distance through time in each individual trial and thick lines the mean over all trials with that number of objects. **(A)** With fixed number of 2 (red), 4 (green), 8 (blue) and 16 (black) objects. **(B)** With increasing number of objects. **(C–D)** Area of convex hull of object positions over time (for $N = 2$ distance is used). Thin lines show the area of the convex hull through time in each individual trial and thick lines the mean over all trials with that number of objects. When calculating the mean over all trials the area of the convex hull of a trial that has finished is set to 0. **(C)** With fixed number of 2 (red), 4 (green), 8 (blue) and 16 (black) objects. **(D)** With increasing number of objects.

completion and standard deviation (time steps) for 2 objects = 68.3 ± 7.3, 4 objects = 130.0 ± 5.6, 8 objects = 273.8 ± 20.2 , and 16 objects = 670.5 ± 108.0. **Figure 3A,C** also confirms that the initial configurations of objects were different in each trial as the initial average object to collection zone distances and convex hulls are different. The relatively low variation in completion times and the fact that initial configurations were different suggests that the process is robust with respect to the initial configurations of objects.

By comparing **Figure 3A,B** (and **Figure 3C,D**) we see that the mean completion time for the case of fixed $N = 16$ and the case with an increasing number of objects are similar. In addition, by comparing the time evolution of the process we see that the process with an increasing number of objects reaches the milestones 2, 4 and 8 objects around the same time that the corresponding fixed number of object trials finishes. This suggests that the process is adaptive with respect to changes in the number of objects, and that potential time and/or efficiency losses associated with its operation in the case of an increasing number of objects versus a fixed number of objects are small.
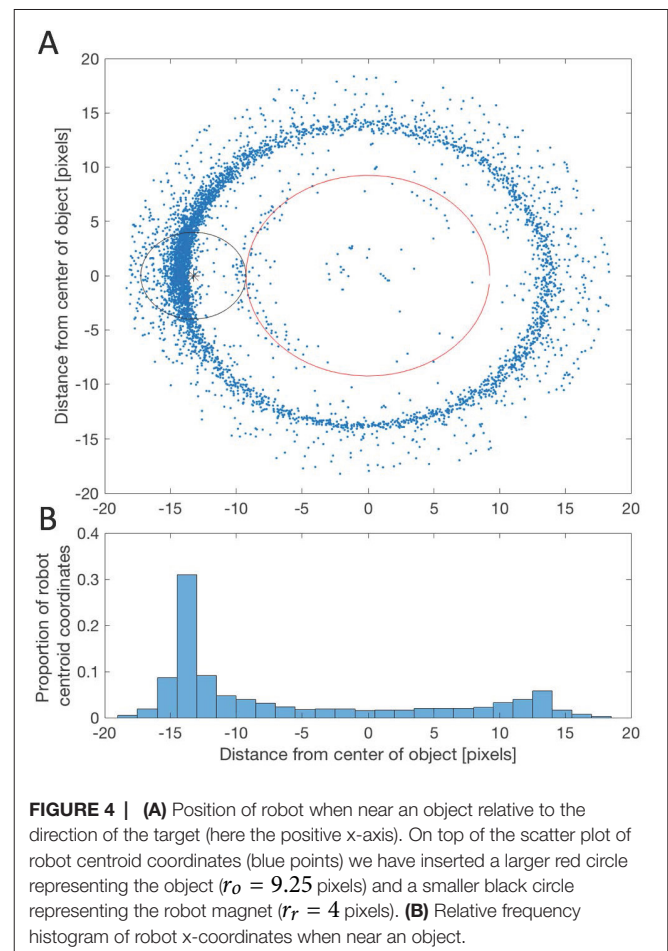
3.2. Robot Behavior

Robot-object interactions are dominated by appropriate collection maneuvers by the robot. When close to an object (within $2r_o$) the robot spends a majority of the time directly behind it relative to the target as presented in **Figure 4A** where the a majority of the robot centroids (blue dots) are on the far side of the object relative to the target. In particular, there is a dense cluster of robot centroids with x-coordinates ranging from about -13 to -15 (**Figure 4A**), which appears to be the ideal position from which to drive the object to the target (**Figure 4B**). Indeed, the peak at +13 to +15 in **Figure 4B** shows that when the robot is on the same side of the object as the target it often pushes the object directly away from the target while attempting to get around it. This is also reflected in the short increases before linear decreases in the measures provided in **Figure 4A,B**. This phenomenon is a consequence of the fact that when the robot is initially approaching an object it often comes directly from the collection zone having just delivered another object. Note that there are some blue dots closer to the object than the object and robot radii should allow, and in some cases even apparently inside the object. These are the result of rare occasions when the robot magnet partially of fully slip up on top of the object. These situations typically sort themselves out quickly and the robot magnet gets off and continues to push the object within a few time steps. However, if the process is supervised inducing a small perturbation to the object or robot can help resolve it even faster.

4. DISCUSSION

We have shown that our biomimetic collection algorithm works when implemented into a simple robot and that the resulting robot collection process exhibits several potentially useful properties.

The collection process is robust with respect to the initial configurations of objects, in the sense that differences in initial

FIGURE 4 | (A) Position of robot when near an object relative to the direction of the target (here the positive x-axis). On top of the scatter plot of robot centroid coordinates (blue points) we have inserted a larger red circle representing the object ($r_o = 9.25$ pixels) and a smaller black circle representing the robot magnet ($r_r = 4$ pixels). **(B)** Relative frequency histogram of robot x-coordinates when near an object.

configuration of objects does not lead to large differences in completion time (**Figure 3A,C**). This result therefore indicates that this process may be a good candidate for reliable collection of objects in novel and noisy environments. In addition, the process is adaptive with respect to changes in the number of objects (**Figure 3A,C**). So it may operate in a changing environment as well as fixed. Finally, there are no obvious time and/or efficiency losses associated with its operation in a changing environment as compared to a fixed environment (comparing **Figure 3A,B**, and **Figure 3C,D**) which would suggest that the cost of operation in a changing environment is effectively the same as in a fixed environment.

We have established that the robot-object interactions are dominated by appropriate collection maneuvers by the robot (**Figure 4**), and that the resulting robot behavior is consistent with the behavior exhibited by sheepdogs and simulated shepherds herding sheep/agents (cf. Figure 5ab, Strömbom et al., 2014). In particular, comparing Figure 5ab from Strömbom et al., 2014 with **Figure 4B** presented in our results, shows that the real dog (Figure 5a, Strömbom et al., 2014), the simulated shepherd (Figure 5b, Strömbom et al., 2014), and the robot (**Figure 4B**) all exhibit distance from the center of flock/object distributions that are skewed with one dominant peak. That the robot-object interactions are dominated by appropriate collection maneuvers by the robot shows that the underlying algorithm and the

implementation into the robot are robust with respect to noise. We know that there are several sources of noise/error in our experimental trials, which are, in order of estimated impact: (i) the robot controlled magnet is not fixed exactly at the center of the robot but has some flexibility, (ii) fluctuations in the time it takes to send instructions to the robot from Matlab via Bluetooth, (iii) image acquisition (any variation in lighting conditions) and centroid calculation error, and (iv) noise in the electrical components (in particular the robot itself). Moreover, whilst our robot does not "behave optimally" (e.g., the robot sometimes pushes gathered objects outside of the collection zone when on route to collect others) its operation is robust and it does, on average, perform well. For our purposes, this is a positive result because it reflects a reality of biological systems, and we did not set out to minimise a cost function (Pérez-Escudero et al., 2009).

Due to the above listed properties of the collection process and its implementation into this simple e-puck robot, in particular its robustness and adaptability, we believe that the algorithm presented here could potentially be used to reliably and effectively collect objects from the environment both on land and on the surface of water if implemented into an appropriate robot. To directly use the implementation presented here, including the feedback control loop, the robot could work as part of a pair, with a surveillance drone that provides the collection robot with overhead images. Considering how accessible advanced drone technology is today this should not present an obstacle. Such a pair consisting of one collection/guiding robot and one surveillance drone could potentially solve a number of problems that are impossible, dangerous, and/or costly for humans to deal with directly. For example, moving animals from sensitive areas (DeVault et al., 2011), removing or limiting the spread of oil on water (Zahugi et al., 2012; Fingas, 2016), collecting hazardous materials (Nguyen et al., 2002), guiding people to safety in areas/rooms with low visibility (Isobe et al., 2004), and potentially even for evacuation and rescue from disaster sites (Patterson et al., 2013).

Finally, we expect that integrating our approach of emphasizing two-way robot-individual interactions into advanced frameworks for animal-robot interactions (e.g. Swain et al. 2012; Bonnet et al. 2017), will afford a greater integration of function and mechanism in the study of collective animal behavior. In particular, it would allow the use of robots to investigate phenomena thought to be intimately linked with specific identifiable individuals, e.g., influential leaders (Jiang et al., 2017). For example, using a robot

with two-way interaction would allow for a precise and dynamic manipulation of leadership traits (played out by a robot) enabling a more standardized, repeatable experimental design and causal analysis of leader-follower dynamics (Nakayama et al., 2012) and their consequences for group-level patterns of behaviour (Cazenille et al., 2017).

AUTHOR CONTRIBUTIONS

AK and DS planned and designed the study, and wrote the paper. DS constructed the arena and feedback control loop, performed the experiments, and processed and analyzed the data.

FUNDING

ACKNOWLEDGEMENT

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/article/10.3389/frobt.2018.00048/full#supplementary-material

VIDEO S1 | Shows one collection trial each for 2, 4, 8 and 16 objects and one trial with an increasing number of objects. All relevant calculated quantities have been superimposed on the overexposed webcam image used as input to the image analysis part. These quantities are: Target (blue asterisk), Collection zone (green ring), Object centroids (red asterisks), Object boundaries (red ring), Robot controlled magnet centroid (black asterisk), Robot controlled magnet boundary (black ring), Current heading (red rod), New (ideal) heading (green rod).

REFERENCES

Bennett, B., and Trafankowski, M. (2012). "A comparative investigation of herding algorithms" *Proc. Symp. on Understanding and Modelling Collective Phenomena (UMoCoP)* 33–38.

Bonnema, G. M. (2012). System design of a litter collecting robot. *Procedia Comput. Sci.* 8, 479–484. doi: 10.1016/j.procs.2012.01.089

Bonnet, F., Cazenille, L., Gribovskiy, A., Halloy, J., and Mondada, F. (2017). "Multi-robot control and tracking framework for bio-hybrid systems with closed-loop interaction," in *Robotics and Automation (ICRA. 2017 IEEE International Conference on* (IEEE) 4449–4456.

Cazenille, L., Collignon, B., Chemtob, Y., Bonnet, F., Gribovskiy, A., Mondada, F., et al. (2017). How mimetic should a robotic fish be to socially integrate into zebrafish groups? *Bioinspir. Biomim.*, in press. doi: 10.1088/1748-3190/aa8f6a

DeVault, T. L., Belant, J. L., Blackwell, B. F., and Seamans, T. W. (2011). Interspecific variation in wildlife hazards to aircraft: Implications for airport wildlife management. *Wildlife Society Bulletin* 35 (4), 394–402. doi: 10.1002/wsb.75

Faria, J. J., Dyer, J. R. G., Clément, R. O., Couzin, I. D., Holt, N., Ward, A. J. W., et al. (2010). A novel method for investigating the collective behaviour of fish: introducing 'Robofish'. *Behav. Ecol. Sociobiol. (Print).* 64 (8), 1211–1218. doi: 10.1007/s00265-010-0988-y

Fingas, M. (2016). *Oil Spill Science and Technology*. Gulf professional publishing.

Gauci, M., Chen, J., Li, W., Dodd, T. J., and Gross, R. (2014). "Clustering objects with robots that do not compute," in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems (International Foundation for Autonomous Agents and Multiagent Systems)* 421–428.

Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., et al. (2007). Social integration of robots into groups of cockroaches to control self-organized choices. *Science* 318 (5853):1155–1158. doi: 10.1126/science.1144259

Hamann, H., Khaluf, Y., Botev, J., Divband Soorati, M., Ferrante, E., Kosak, O., et al. (2016). Hybrid societies: challenges and perspectives in the design of collective behavior in self-organizing systems. *Front. Robot. AI* 3. doi: 10.3389/frobt.2016.00014

Hamilton, W. D. (1971). Geometry for the selfish herd. *J. Theor. Biol.* 31 (2), 295–311. doi: 10.1016/0022-5193(71)90189-5

Handegard, N. O., Boswell, K. M., Ioannou, C. C., Leblanc, S. P., Tjøstheim, D. B., and Couzin, I. D. (2012). The dynamics of coordinated group hunting and collective information transfer among schooling prey. *Curr. Biol.* 22 (13), 1213–1217. doi: 10.1016/j.cub.2012.04.050

Hubert, J., and Weibel, Y. (2008). *ePic2 v2.1 Documentation* . https://www.joomlashine.com/documentation/jsn-templates/jsn-epic/jsn-epic-2-configuration-manual.html.

Ioannou, C. C., Guttal, V., and Couzin, I. D. (2012). Predatory fish select for coordinated collective motion in virtual prey. *Science* 337 (6099), 1212–1215. doi: 10.1126/science.1218919

Isobe, M., Helbing, D., and Nagatani, T. (2004). Experiment, theory, and simulation of the evacuation of a room without visibility. *Phys. Rev. E Stat. Nonlin. Soft. Matter. Phys.* 69 (6 Pt 2):066132. doi: 10.1103/PhysRevE.69.066132

Jiang, L., Giuggioli, L., Perna, A., Escobedo, R., Lecheval, V., Sire, C., et al. (2017). Identifying influential neighbors in animal flocking. *PLoS Comput. Biol.* 13 (11):e1005822. doi: 10.1371/journal.pcbi.1005822

Karunasena, G., Amaratunga, R., and Haigh, R. (2008). "Automation and robotics in post disaster waste management: post tsunami Sri Lanka" *In 25th International Symposium on Automation and Robotics in Construction*

King, A. J., Wilson, A. M., Wilshin, S. D., Lowe, J., Haddadi, H., Hailes, S., et al. (2012). Selfish-herd behaviour of sheep under threat. *Curr. Biol.* 22 (14), R561–2. doi: 10.1016/j.cub.2012.05.008

Krause, J., Winfield, A. F., and Deneubourg, J. L. (2011). Interactive robots in experimental biology. *Trends Ecol. Evol. (Amst.).* 26 (7), 369–375. doi: 10.1016/j.tree.2011.03.015

Landgraf, T., Bierbach, D., Nguyen, H., Muggelberg, N., Romanczuk, P., and Krause, J. (2016). RoboFish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live Trinidadian guppies. *Bioinspir. Biomim.* 11 (1):015001. doi: 10.1088/1748-3190/11/1/015001

Landgraf, T., Nguyen, H., Forgo, S., Schneider, J., Schröer, J., and Krüger, C. (2013). "Interactive robotic fish for the analysis of swarm behavior," in *International Conference in Swarm Intelligence*. Springer 1–10.

Lien, J. -M., Bayazit, O. B., Sowell, R. T., Rodriguez, S., and Amato, N. M. (2004). "Shepherding behaviors," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on (IEEE)*, Vol. 4 4159–4164.

Lien, J-M., Rodriguez, S., Malric, J-P., and Amato, NM. (2005). "Shepherding behaviors with multiple shepherds" *In Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on (IEEE)* 3402–3407.

Melhuish, C., Wilson, M., and Sendova-Franks, A. (2001). "Patch sorting: Multi-object clustering using minimalist robots," in *European Conference on Artificial Life*. Springer 543–552.

Miki, T., and Nakamura, T. (2006). "An effective simple shepherding algorithm suitable for implementation to a multi-mmobile robot system," in *Innovative Computing, Information and Control, 2006. ICICIC'06. First International Conference on (IEEE)*, Vol. 3 161–165.

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., and Klaptocz, A. (2009). "The e-puck, a robot designed for education in engineering" *Proceedings of the 9th conference on autonomous robot systems and competitions* (IPCB: Instituto Politécnico de Castelo Branco), 59–65.

Nakayama, S., Harcourt, J. L., Johnstone, R. A., and Manica, A. (2012). Initiative, personality and leadership in pairs of foraging fish. *PLoS ONE* 7 (5):e36606. doi: 10.1371/journal.pone.0036606

Nguyen, T. N., O'Donnell, C., and Nguyen, T. B. (2002). "Multiple autonomous robots for uxo clearance, the basic uxo gathering system (bugs) project," in *Multi-Robot Systems: from Swarms to Intelligent Automata*. Springer 53–61.

Patterson, MC., Mulligan, A., and Boiteux, F. (2013). "Safety and security applications for micro-unmanned surface vessels" *Oceans-San Diego, 2013 (IEEE)* (San Diego, CA, USA), 1–6.

Pérez-Escudero, A., Rivera-Alba, M., and de Polavieja, G. G. (2009). Structure of deviations from optimality in biological systems. *Proc. Natl. Acad. Sci. U.S.A.* 106 (48), 20544–20549. doi: 10.1073/pnas.0905336106

Pérez-Escudero, A., Vicente-Page, J., Hinz, R. C., Arganda, S., and de Polavieja, G. G. (2014). idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nat. Methods* 11 (7):743–748. doi: 10.1038/nmeth.2994

Shi, Q., Ishii, H., Kinoshita, S., Takanishi, A., Okabayashi, S., Iida, N., et al. (2013). Modulation of rat behaviour by using a rat-like robot. *Bioinspir. Biomim.* 8 (4):046002. doi: 10.1088/1748-3182/8/4/046002

Strömbom, D., Mann, R. P., Wilson, A. M., Hailes, S., Morton, A. J., Sumpter, D. J., et al. (2014). Solving the shepherding problem: heuristics for herding autonomous, interacting agents. *J. R. Soc. Interface* 11 (100):20140719. doi: 10.1098/rsif.2014.0719

Swain, D. T., Couzin, I. D., Ehrich Leonard, N., and Leonard, N. E. (2012). Real-time feedback-controlled robotic fish for behavioral experiments with fish schools. *Proc. IEEE* 100 (1), 150–163. doi: 10.1109/JPROC.2011.2165449

Turgut, A. E., Çelikkanat, H., Gökçe, F., and Şahin, E. (2008). "Self-organized flocking with a mobile robot swarm," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, Vol. 2. US: Springer, 97–120. doi: 10.1007/s11721-008-0016-2

Vaughan, R., Sumpter, N., Frost, A., and Cameron, S. (1998). "Robot sheepdog project achieves automatic flock control," in *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour.* 489: 493.

Zahugi, E. M., Shanta, M. M., and Prasad, T. (2012). Design of multi-robot system for cleaning up marine oil spill. *Int. J. Adv. Info. Tech.* 2 (4), 33–43. doi: 10.5121/ijait.2012.2404

Cooperative Object Transport in Multi-Robot Systems: A Review of the State-of-the-Art

Elio Tuci[1]*, Muhanad H. M. Alkilabi[2] and Otar Akanyeti[3]

[1] *The Department of Computer Science, Middlesex University, London, United Kingdom,* [2] *The Department of Computer Science, University of Kerbala, Karbala, Iraq,* [3] *The Department of Computer Science, Aberystwyth University, Aberystwyth, United Kingdom*

In recent years, there has been a growing interest in designing multi-robot systems (hereafter MRSs) to provide cost effective, fault-tolerant and reliable solutions to a variety of automated applications. Here, we review recent advancements in MRSs specifically designed for cooperative object transport, which requires the members of MRSs to coordinate their actions to transport objects from a starting position to a final destination. To achieve cooperative object transport, a wide range of transport, coordination and control strategies have been proposed. Our goal is to provide a comprehensive summary for this relatively heterogeneous and fast-growing body of scientific literature. While distilling the information, we purposefully avoid using hierarchical dichotomies, which have been traditionally used in the field of MRSs. Instead, we employ a coarse-grain approach by classifying each study based on the transport strategy used; pushing-only, grasping and caging. We identify key design constraints that may be shared among these studies despite considerable differences in their design methods. In the end, we discuss several open challenges and possible directions for future work to improve the performance of the current MRSs. Overall, we hope to increasethe visibility and accessibility of the excellent studies in the field and provide a framework that helps the reader to navigate through them more effectively.

Keywords: multi-robot systems, cooperative object transport, pushing, pulling, caging

INTRODUCTION

This paper reviews recent research works in MRSs targeting cooperative object transport scenario. A MRS is robotic system consisting of more than one robot (see Cao et al., 1997). MRSs are a promising alternative to automate tasks that are beyond the competency of single robot systems. Transporting big objects, surveillance of vast areas, or robot tasks that can be decomposed into smaller tasks so that they can be carried out simultaneously by several robots are examples of application domains particularly suited for MRSs (Yan et al., 2013). In addition, MRSs, comprised of many but simple individuals, may be cheaper to build and easier to program than a complex robot capable of performing similar tasks (Farinelli et al., 2004; Cai and Yang, 2012; Yan et al., 2013; Khamis et al., 2015; Jiang et al., 2016). MRSs are also potentially more resilient to a large variety of hardware or software failures; when one robot fails or makes a mistake, the others can still complete the task successfully (Parker, 1998).

Although the members of a MRS can be designed or programmed to compete with each other (see Martín H. et al., 2010), the majority of the previous studies have investigated how group members

can work together to achieve a common goal (i.e., cooperation). However, so far the scientific community has failed to agree on a formal definition for cooperation. For some authors, it is sufficient to refer a MRS as cooperative as long as its members share a common goal, even if they have zero interaction (Wang et al., 1994; Quinn, 2004). For others, the definition of cooperation is more strict. A MRS is assumed to be cooperative only if the robot task can not be serialised (i.e., single robot can not complete the task in a sequential manner), and specific cooperation mechanisms should be in place so that the robots can coordinate their actions, and possibly complement each others' capabilities (see Kube et al., 1993; Brown and Jennings, 1995; Cao et al., 1997; Iocchi et al., 2000; Yan et al., 2013). The underlying process that enables cooperative MRSs is generally referred to as coordination of actions (see Kube and Bonabeau, 2000; Simmons et al., 2001; Emery et al., 2002; Farinelli et al., 2004).

Here, for the first time, we provide a comprehensive review on research studies that focus on one application domain; cooperative object transport, the term is coined after (Groß and Dorigo, 2004). Cooperative MRSs are generally employed when the object is too heavy, too large, or has a complex shape so that it can not be transported by a single robot. However, this is not a strict requirement; not all group members need to participate in the physical act of transport; carrying or pushing/pulling the object. Cooperation can still be achieved when a single or few robots transport the object, and the others plan the coordination and navigation of the transporters along a desired trajectory, or clear the way from obstacles (e.g., see Habibi et al., 2015).

Autonomous MRSs capable of cooperative object transport can be extremely effective in a variety of applications that have high economic and societal impact potential; e.g., waste retrieval and disposal, de-mining, or operations requiring object manipulation in environments where direct human intervention is impossible or impractical, such as in space or in deep sea (Huntsberger et al., 2000; Parker and Zhang, 2006; Woern et al., 2006). Thanks to parallelism and decentralised nature of MRSs, the robots apply spatially distributed forces (i.e., pushing, pulling or lifting at different locations) around objects. The physical separation and the independent actions of different agents can potentially generate a group dexterity that a single robot can hardly achieve, irrespective of its sophistication and power (see Brown and Jennings, 1995). This property is particularly important in cooperative transport tasks, where the independent exertion of multiple pushing/pulling forces in different points of an object can allow the group to generate precise translation/rotation manoeuvres in order to avoid obstacles during transport.

Due to its relevance, cooperative transport has been studied in recent years by research works that have extensively looked at different aspects related to the coordination and synchronisation of the forces required to initiate and sustain the transport of objects that can not be transported by a single robot. The research on cooperative transport in MRSs has been progressing by investigating and testing the potentialities of a variety of different methodological approaches, that are generated by integrating, with different modalities, the various available alternatives for what concerns methods and techniques to design the mechanisms underpinning the desired group responses, means for inter-robot

communications, transport techniques, evaluation scenarios, etc. The objective of this paper is to review and at the same time to provide a navigation framework to order and critically evaluate this rather heterogeneous and fast growing body of literature. We employ a rather coarse-grain categorisation system that distinguishes and orders the research works with respect to the type of transport strategy used by the group to cooperatively move the object. We believe that this categorisation system represents a helpful perspective to account for the scientific progress made by a methodologically diverse body of literature, and to identify open challenges and promising directions for further work to improve the transport capabilities of MRSs.

We review and categorise the research works using three categories, each of which is discussed in a separate section:

1. Pushing-only strategy (see section 2): robots are not physically attached to the object, and transport is achieved by pushing the object.
2. Grasping strategy (see section 3): robots are physically attached to the object, and transport is achieved by either pushing or pulling (or both) the object.
3. Caging strategy (see section 4): this strategy is similar to the pushing-only strategy. Robots are distributed to entrap the object (i.e., caging) and they hold the object tightly during transport.

We decided to separate pushing-only and caging strategies even though they share some characteristics. This is because the latter is not only concerned with transporting the object but also maintaining an object closure at all times. This additional requirement imposes unique design challenges which influence the communication and coordination strategies employed by the robots (see Hekmatfar et al., 2014). The reader should be aware that cooperative transport has also been studied in MRSs that, due to their characteristics, they do not fit in any of our three categories. In particular, cooperative transport has been studied in a group of aerial robots (Michael et al., 2011; Bernard et al., 2011) required to carry heavy objects using cables. Moreover, cooperative transport has gained significant attention in micro-scale applications where micro-robots (i.e., robots with sub-millimetre or smaller dimensions capable of manipulating micro-objects including living cells) have been designed to perform micro-manipulation and micro-assembly tasks such as molecular delivery to targeted cells, minimally invasive surgeries, tissue engineering, and other general micro-manipulation applications (Hu et al., 2011; Shahrokhi and Becker, 2016; Rahman et al., 2017). We decided to exclude from this review these and other similar research works based on transport strategies alternative to the pushing, grasping, and caging strategy described above.

In section 5, we provide an informative and constructive discussion on the state of the art of MRSs engaged in cooperative transport that helps to identify objectives for interesting future directions of research. Contrary to other similar review papers, we do not employ the classic and frequently used dichotomous view that distinguishes MRSs in those controlled in a centralised and those controlled in a decentralised way (see Cao et al., 1997; Bahçeci et al., 2003; Bayindir and Şahin, 2007). We believe that, in the context of cooperative transport, the use of such a dichotomous

perspective would blur important methodological details that largely contribute to the identity and the originality of every single study. We rather complement the review framework based on the type of transport strategy illustrated above, with references to the eventual presence of any *key element* in the MRS's architecture, and we comment on the type of communication used to achieve the coordination of action among the group members. In our view, the *key element* can be either a member of the group or an element external to the group (e.g., a server), that orchestrates the dynamics of the group by regulating some or the totality of the actions of those agents that are subordinated to its decisions. When the *key element* is internal to the group, it is generally represented by an agent that is either structurally or functionally different from the other robots of the group (e.g., a leader). When the competencies and contributions of the *key element* to the group performance can not be dynamically allocated and more importantly re-allocated to any of the other members of the group, the *key element* undoubtedly represents a single-point of failure of the system. This is because a failure of the *key element* inevitably leads to the failure of the entire system.

Before we start, we would like to clarify few points. First, our review mainly focuses on research studies that use mobile robots. These robots typically vary in body length, and methods for locomotion (e.g., legged or wheeled robots). The studies using other types of robots (e.g., aerial and aquatic robots, and micro/macro robotic manipulators) are omitted unless there is a specific point to be made. Second, even within the mobile robotics literature, there is a large body of work, which is impossible to cover in one review paper. Hence, we try to select the most representative studies that employ different transport, coordination and control strategies. Third, we define four terms that help us to better describe various MRSs approaches: MRSs that use direct or indirect communication, and homogeneous and heterogeneous MRSs. In direct communication, the members of a MRS send/receive messages to/from each other using a dedicated communication network. Messages are often transmitted via text, sound or light using wireless communication protocols. Based on these protocols, message exchange can be private (i.e., between two or selected group members), local (i.e., among neighbours within close proximity) or global (among all members). In indirect communication, the robots are not allowed to communicate with each other explicitly. Instead, they communicate implicitly using the object they transport and/or through the changes in the environment they operate in. In homogeneous MRSs, all group members are identical with same hardware (i.e., physical) and software (i.e., functional) designs, whereas in heterogeneous MRSs, at least one group member is physically and/or functionally different from the others. Homogeneous groups are more frequently found in swarm robotics, a sub-field of the MRSs research area where the robots mimic main characteristics and behaviour of social insects, such as ants and bees (see Şahin, 2005). The structural and functional homogeneity of a robotic swarm is inspired to the genetic similarities of "workers" in social insects. The group homogeneity is supposed to make the group more scalable with respect to its size and more resilient to individual failure, since in principle any robot can replace any other identical member of the group. The group homogeneity does not preclude the possibility that a certain

amount of functional diversification could characterise the group members, as long as any behavioural specialisation emerges during the life of the group, and it is in principle reversible. Cooperative object transport scenarios often require complex and diversified behavioural competencies that scientists have very frequently implemented by exploiting structurally and/or functionally heterogeneous rather than homogeneous groups. Advantages and drawbacks of the use of heterogeneous groups in the context of cooperative transport will be further discuss in section 5.

THE PUSHING-ONLY STRATEGIES

Pushing-only strategies are methods of collectively transporting items by exerting pushing forces on the item. These type of strategies are primarily employed by robots that can not pull objects, since they have no means to grasp them. Pushing-only strategies may appear to be relatively simple methods of cooperative transport. However, on top of the challenges common to all transport strategies (e.g., the alignment of forces required to initiate the transport, etc.), pushing-only strategies require a significant amount of coordination of actions to sustain the transport. The item may move on a very inefficient trajectory unless the robots carefully manage frictional, gravitational, and dynamical forces to stabilise the direction of transport. **Table 1** summaries the main characteristics of the research works reviewed in this section. Generally speaking, it is worth noticing that the large majority of these works are based on homogeneous groups, where the robots' controller is designed using a behaviour-based methodology (see Brambilla et al., 2013, for further details). Groups exploiting indirect communication prevail on groups exploiting forms of direct communication. Half of the studies look at a simplified transport scenarios, where the problems related to the initial alignment of pushing forces is solved by initialising the robots very close to the object, facing the same side of the object (see no random initial positions in **Table 1**). In the following, we review these works, by emphasising objectives and achievements.

The study in (Kube et al., 1993) can be considered the pioneering work targeting a cooperative transport task by a homogeneous group of simple robots that can only push the object (i.e., a box). This study is considered to be the first research work that formally represented in "hardware" the dynamics of cooperative transport. The authors demonstrate that coordinated efforts in a box pushing task are possible without the use of direct communication or robot differentiation. The group exploits the physical interactions among the robots and between the robots and the object to initiate and to sustain the transport. In (Kube and Zhang, 1997; Kube and Bonabeau, 2000), the authors further develop the model described in (Kube et al., 1993) with the addition of a stagnation recovery strategy. Stagnation refers to a deadlock condition in which the robots cancel each others' pushing forces due to the way in which they are positioned around the object. The authors also evaluate the group transport strategies with objects of different shapes in scenarios in which the objects have to be transported towards a moving target.

Mataric et al. (1995) propose the use of direct communication to improve the coordination of a homogeneous group of two six-

TABLE 1 | Main characteristics of research works using a pushing-only transport strategy.

	Experimental setting					Robustness Test			Control and communication			Similar studies
	nr. robots	group type: -homogeneous (ho) -heterogeneous (he)	env. with obstacles	obj. shape	transport direction	random initial positions	obj. mass	obj. dim	controller design method	autonomous allocation of roles	communication	
Kube et al., 1993	5	ho	no	rectangular prism	arbitrary	no	no	no	behaviour based	n/a	indirect	Kube and Zhang, 1997; Kube and Bonabeau, 2000
Mataric et al., 1995	2	ho	no	rectangular prism	towards a goal	no	no	no	behaviour based	n/a	direct	
Gerkey and Matarić (2002)	3	he	no	rectangular prism	towards a goal	no	no	no	other	Yes	direct	
Yamada and Saito, 2001	4	ho	no	rectangular prism	towards a goal	no	no	no	behaviour based	n/a	indirect	
Jianing Chen et al., 2015	20	ho	no	cylinder, rectangular, triangular prism	towards a goal	yes	yes	no	behaviour based	n/a	indirect	Kapellmann-Zafra et al., 2016
Sugie et al., 1995	4	ho	no	rectangular prism	towards a goal	yes	no	no	other	n/a	direct	
Wang and de Silva, 2006a	2	he	yes	rectangular prism	towards a goal	yes	no	no	automatic	no	direct	
Alkilabi et al., 2017	2 to 6	ho	no	rectangular prism	arbitrary	yes	yes	yes	automatic	n/a	indirect	Alkilabi et al., 2016
Fujisawa et al., 2013	4 to 10	ho	no	cylinder	towards a goal	yes	no	no	behaviour based	n/a	indirect	
Neumann et al., 2014	2	ho	no	rectangular prism	predefined path	yes	no	no	other	n/a	indirect	
Tang and Parker, 2005	2 to 3	he	no	rectangular prism	arbitrary	Yes	no	no	behaviour based	no	direct	
Parker, 2000	2	he	no	rectangular prism	arbitrary	no	no	no	other	no	direct	
Wang and de Silva, 2006b	3	ho	yes	rectangular prism	towards a goal	no	no	no	automatic	n/a	indirect	Sen et al., 1994

A transport direction is classified as "arbitrary" when the robots autonomously choose where to transport the object. The "no random initial positions" category refers to specific initial spatial arrangements of the robots relative to the object that tend to facilitate the alignment of pushing forces. Robustness tests with respect to the object mass and dimensions are tests in which the effectiveness of group transport strategies are evaluated with objects of different mass, and different size. The classification of the controller design methods is the one discussed in (Brambilla et al., 2013).

legged robots required to cooperatively transport a rectangular box toward a target. Published during a time of disaffection for the classic AI paradigm, this study aims to demonstrate that tasks requiring complex coordination of actions among physical robots can be successfully accomplished without the robots having any model of the world and without being able to make any predictions on the consequences of their actions. Robots' controller is designed using a behaviour based methodology (Brooks, 1986), and communication is used by the agents to exchange their sensors readings and to implement a turn-taking protocol. To facilitate the initial alignment of pushing forces, the robots are positioned on the left and on the right end of one of the longest object's side. The results indicate that the use of communication and of the turn-taking protocol significantly helps the robots to improve the overall group performance.

Gerkey and Matarić (2002) illustrate the performances of a group of three robots in which one element of the group plays the role of the watcher, and the other two robots play the role of the pusher. The watcher perceives the object and the goal destinations, and its main duty is to lead the team by providing the other robots information concerning the direction of transport. The pushers push the cuboid object without perceiving the goal destination which remains hidden behind the box that occludes their view. The robots rely on a direct form of communication for the coordination of their actions. The transport trajectory is free from obstacles, and roles are assigned using an auction-based system (i.e., MURDOCH architecture, see Gerkey and Matarić, 2001). The heterogeneous group manages to successfully transport the object in straight and curved trajectories. The system also proved to be resilient to the failure of one of the pusher, and to a certain extent to the failure of the communication mechanism underpinning the watcher-pusher interaction. However, the system heavily relies on the capabilities of the single watcher, which acts as a *key element* that gathers sensory information sent by the pushers and generate the group response by instructing the pushers on how to move.

The study illustrated in (Yamada and Saito, 2001) is also conceived in support of a theoretical perspective alternative to the classic AI, since its main goal is to demonstrate with physical robot experiments that an environment selection task and a cooperative box-pushing task can be both carried out by a homogeneous group of robots where agents are guided by a reactive controller. Contrary to (Mataric et al., 1995) and (Gerkey and Matarić, 2002) which advocate for the use of direct communication, the results illustrated in (Yamada and Saito, 2001) demonstrate that indirect communication is sufficient to cooperatively transport an object toward a target area. The robots can operate in a simple environment where individual robots are required to push light boxes, or in complex environments where multiple robots are required to cooperatively push a heavy box. The mechanisms underlying the environment selection task operate under the assumption that there is no moving object except the robots. Moreover, it is assumed that during pushing, no wheels slippage is experienced by the robots even in those cases in which the object does not move when subject to pushing forces. These assumptions are required to allow the robots to discriminate between those cases in which the box is light enough to be transported individually, from

those cases in which the box is so heavy to require a cooperative response.

Jianing Chen et al. (2015) propose an alternative group transport method which exploits occlusion, rather than trying to overcome the limitations imposed by it. The robots are designed to push the object across the portion of its surface, where it occludes the direct line of sight to the goal. In this study, a group of twenty e-puck robots (see Mondada et al., 2009) are required to transport a cylindrical object towards a goal. The robots push the object only when they can not see the goal destination. This simple behaviour results in transporting the object towards the goal without using any form of direct communication. The authors also provide an analytical proof of the effectiveness of the method, and results of successful empirical tests with a cuboid and a triangular objects are discussed. In (Kapellmann-Zafra et al., 2016), the occlusion-based strategy discussed in (Jianing Chen et al., 2015) is tested in a task in which the robots are required to transport an object towards a moving target, represented by another robot.

The study described in (Sugie et al., 1995) is one of the first to address the problem of designing push-only strategies in a dynamic environment that incorporates obstacles. The authors describe a system in which the robots infer other robots' intentions by observing their behaviour and cooperate based on those inferred intentions. A camera placed on the ceiling of the robots arena communicate to each robot the position of all other robots, obstacles, boxes to be transported, and final destinationsof each box. An algorithm made of a task planner, a pushing action planner, and a dynamic obstacle avoidance function guides the robots during the task execution. In this as in other similar studies in which the control algorithm relies on a global view of the environment, the group transport strategy, although particularly effective to manoeuvre the object in a complex environment with obstacles, would not tolerate a fault to its multiples *key elements*, such as the camera and the task planner.

Wang and de Silva (2006a) consider a heterogeneous group of robots that is required to cooperatively transport a box by removing obstacles that abstract the way to the final transport destination. The authors propose an approach based on the use of a force/motion control system. Three different types of agent are used in this approach: a vision agent that has a global view of the environment to generate positions and orientation coordinates of all robots, the object, and the obstacles; a learning agent responsible for generating cooperation plans based on an optimisation approach that integrates reinforcement learning and genetic algorithm; two physical robots that execute the plan generated by the learning agents. The plan may require one robot to leave the transport to remove obstacle/s obstructing the way to the final transport destination. The study demonstrates the feasibility and the effectiveness of the proposed method using experiments with two small prototype robots. Both the vision and the learning agents are *key elements* whose contribution is vital for the correct functioning of the MRS.

Alkilabi et al. (2017) demonstrate that effective coordination of actions for initiating and sustaining the transport of heavy objects to be moved in an arbitrary direction can be obtained by homogeneous groups of robots by exploiting a relatively simple form of indirect communication based only on the possibility to perceive the movements of the object. In this study, physical e-puck

robots are equipped with an optic-flow sensor whose readings are used to distinguish between cases in which the robots pushing forces contribute to moving the object from those cases in which the robots efforts do not result in any significant object movement. The possibility to discriminate between the above mentioned two circumstances is vital for the initial alignment of pushing forces and for sustaining the transport. The authors show that the transport strategies are scalable with respect to the group size, and robust enough to deal with boxes of various mass and size. In a complementary study illustrated in (Alkilabi et al., 2016), the authors complement the robots' neuro-controller, initially designed to support the object transport in an arbitrary direction, with mechanisms to direct the transport towards a specific target location.

A cooperative transport study that uses indirect communication via artificial pheromone is described in (Fujisawa et al., 2013). In this study, a group of ten robots can sense and lay on the terrain a volatile alcohol substance that mimics the effect of ants' pheromone during trail formation. The task requires the robots to perform a random search to find a food item (i.e., heavy object), and to transport it to a goal location (i.e., the nest). The pheromone-based communication is used by the robots to recruit other nest-mates when a food location is identified. The results indicate that pheromone-based communication contributes to reducing the task completion time, in comparison to the case in which the robots depend completely on a random walk to congregate at the food. The study also shows that the pheromone-based communication is effective only with a relatively small number of robots in the environment. When a larger group of robots is used, the pheromone-based communication has less impact on the completion time, as many robots are likely to find the food and begin the cooperative transport before the trail is formed.

In (Neumann et al., 2014), an algorithm running on an external server controls a group of robots required to push a box on straight and circular trajectories defined by the experimenter. The algorithm generates informations concerning where the robots have to apply pushing forces and the magnitude of the forces needed to transport the box. Position and orientation of the robots and of the box are measured using an ultra-wide band tags placed on the robots as well as the box. The readings generated by the force sensors and data relative to the robots' position generated by the ultra-wide band system are routed to a central server, which in turn calculates the robots' required speed and sends the commands to the robots accordingly. The robots execute the commands to generate the desired forces and torques on the object in order to move it along a planned trajectory. The study demonstrates the validity of the proposed method using two Pioneer robots equipped with hinged force sensors extension. The server running the control algorithm is the *key element* which manages the robots' actions by sending instructions to each robot using a direct form of communication, supported by a wireless communication network. Such type of direct communication tends to suffer from scalability issue, since the communication load increases when the number of robots increases. This may cause a decrease in system performance or in extreme cases, it can result in an overall system failure. Moreover, the scalability of the transport strategies may be also hindered by issues related

to the design of network topologies and to the communications protocols (see Cao et al., 1997).

The cooperative object transport scenario using a pushing only strategy has also been used in various research studies as a benchmark task to evaluate the functional characteristics of various control policies (see Sen et al., 1994; Parker, 2000; Tang and Parker, 2005; Wang and de Silva, 2006b).

THE GRASPING STRATEGIES

Grasping strategies are methods by which the robots physically attach to an item to be able to collectively transport it. Thus, grasping strategies can only be exploited by robots which possess the mechanisms to grasp an object. There exists a variety of mechanisms that allows a robot to physically connect to an object, some of which allow the robots not only to grasp but also to lift an item. Compared to pushing-only strategies, grasping strategies provide a better control over the transported object, since once grasped, the object can be either pushed or pulled. However, stable and effective grasping strategies often require the robots to optimally distribute around an object in order to avoid undesired effects, such as the object touching the ground, or the load being distributed in an unbalanced way among the robots. To avoid the challenges related to the effective positioning of the robots around the object, the majority of the research works reported in the literature focus on the development of grasping strategies by groups of robots that are pre-attached to the object and optimally positioned around it before starting the transport (see also **Table 2**). The work described in (see Sasaki et al., 1995) is one of the few in which the authors develop an algorithm to allow a homogeneous group of robots to find the optimal arrangement around an object that has to be lifted and transported to a final destination. In this study, the robots know the shape of the object. They estimate the object mass and mass centre position by lifting the object, and they use these estimates to optimally distribute the grasping points around the object.

Most of the research works on cooperative transport using grasping strategies rely on the presence of a robot leader to generate the desired motion trajectory of the object. In these studies, no mechanisms for a dynamic allocation of roles are contemplated. Thus, the leader can be considered a *key element* which, if it fails, the entire group stops working. A leader/follower approach is described in (Kosuge and Oosumi, 1996), where a group of two robots cooperatively transport a long cuboid object pre-attached to them using free rotational joints. The control algorithm requires the presence of a leader robot that is in charge of implementing a specific motion trajectory. The follower supports the leader in the transport of the object along the desired trajectory by coordinating its actions through the perception of the forces applied to the object. In (Kosuge et al., 1998), the authors extend this algorithm originally designed and tested on holonomic robots with velocity-controlled actuators to nonholonomic mobile robots driven by two wheels. In (Takeda et al., 2002), the authors further improve the control algorithm by adding a collision avoidance unit to enable the robots to transport a single object in more complex environments with obstacles.

TABLE 2 | Main characteristics of research works using a grasping transport strategy.

	Experimental setting						Control and communication			Similar studies
	nr. robots	group type: -homogeneous (ho) -heterogeneous (he)	autonomous allocation of roles	env. with obstacles	obj. shape	transport direction	pre-attached	controller design method	communication	
Sasaki et al., 1995	3	ho	n/a	no	irregular convex shaped	object only lifted, not transported	no	other	direct	
Kosuge and Oosumi, 1996	2	he	no	no	rectangular prim	arbitrary	yes	other	direct	Kosuge et al., 1998; Takeda et al., 2002
Wang and Schwager, 2016	4	he	no	yes	rectangular prism	know to leader	yes	other	indirect	
Wang et al., 2016	4	he	no	yes	square prism	known to leader	yes	other	indirect	
Farivarnejad et al., 2016	4	ho	n/a	no	rectangular prim	predefined path known to all robots	yes	other	indirect	
Habibi et al., 2015	5 to 7	he	no	yes	irregular convex shaped	predefined path known to guide robots	yes	other	direct	Habibi et al., 2014
Yufka and Ozkan, 2015	2–3	ho	n/a	no	rectangular prism	towards a goal known to all robots	yes	other	direct	
Tuci et al., 2006	6	ho	n/a	no	cylinder	towards a goal visible by all robots	no	automatic	indirect	Groß and Dorigo, 2008; Gross and Dorigo, 2009
Ferrante et al., 2010	3–4 (sim)	ho	n/a	yes	cylinder	towards a goal known to some robots	yes	behaviour based	direct	Campo et al., 2006
Berman et al., 2011	4–10 (sim)	ho	/ao	no	cylinder	towards a goal known to all robots	no	behaviour based	indirect	Wilson et al., 2014; Gelblum et al., 2016; Guo et al., 2017
Stilwell and Bay, 1993	11 (sim)	he	no	no	object on pallet	towards a goal known to all robots	yes	other	direct	Johnson and Bay, 1995; Bay, 1995; Pereira et al., 2002; Loh and Traechtler, 2012
Hichri et al., 2016	3 (sim)	ho	n/a	no	rectangular prism	towards a goal known to all robots	no	other	direct	Wang et al., 1994; Yamashita et al., 2003
Stroupe et al., 2005	3	he	no	no	rectangular prism	towards a goal known to the robot leader	no	behaviour based	indirect	Hashimoto et al., 1993
Machado et al., 2016	2	he	no	yes	rectangular prism	towards a goal known to the robot leader	yes	behaviour based	indirect	Soares et al., 2007

"pre-attached" refers to an initial configuration in which the robots are already attached to the object. See also caption **Table 1** for further illustration of the descriptive categories.

A leader/follower approach is also exploited in (Wang and Schwager, 2016) and in (Wang et al., 2016). Wang and Schwager (2016) describe a kinematic controller for a group of four robots, in which the robot leader pulls the object and defines the direction of transport, and the robots follower push the object to sustain the leader effort. The model requires the robots to have information beforehand of the friction coefficients, the mass of the object, and the total number of the robots forming the group, in order to measure the velocity and acceleration at the centre of mass of the object. The robots are manually attached to the object with a fixed connection established by a one DOF gripper. Three experimental set-ups are studied with different types of leaders (i.e., an autonomous robot, a robot teleoperated by a human, and a human leader), while the characteristics of the robots follower are kept unchanged in all three experimental set-ups. The results of the study demonstrate the followers' ability to effectively coordinate with all types of leader by following the leader-defined direction of object motion. In (Wang et al., 2016), the kinematic control described in (Wang and Schwager, 2016) is extended in order to allow a group of four custom-built omnidirectional robots (i.e., OuijaBots) to transport a longitudinal object along trajectories requiring the object to be rotated in order to cross a narrow corridor.

The objective of the study described in (Farivarnejad et al., 2016) is to design controllers that drive a homogeneous group of four "Pheeno" robots (see Wilson et al., 2016) to collectively transport a rectangular load at a desired speed along a straight path in a target direction. No distinction in leader/follower is assumed. Moreover, the robots do not have global localisation or communication capabilities, and they lack information about the payload dynamics, the number of robots in the transport team, their distribution around the payload, and the layout of the environment. It is assumed that each robot can measure its speed and heading, and it is given access to the desired target direction of the transport. The position and orientation of the robots with respect to the object are also known since all robots are rigidly pre-attached to the object. Each robot is equipped with wheel encoders to estimate its velocity and a compass to calculate its heading. The results demonstrate the robots' capabilities in transporting the object in relatively straight trajectories parallel to the desired path with some drift caused by the noise in the compass measurements and the errors in the odometry due to the wheel slippage.

In (Machado et al., 2016) and in (Soares et al., 2007) robots are controlled with a dynamic control architecture that uses the attractor dynamic approach to behaviour-based robotics (see Bicho and Schöner, 1997). In the most recent work (see Machado et al., 2016), the authors test the control architecture on a group of two physical robots jointly transporting a rectangular prism carried on a payload support base capable of returning bearing and displacement of the load with respect to the robots centre of mass. The leader robot, equipped with an omnidirectional camera, generates the transport trajectories in order to avoid static and moving obstacles that obstruct the transport. The results of the study show that the dynamic control architecture allows the heterogeneous group of mobile robots to operate in complex cluttered environments and to successfully transport loads of different with and length.

Habibi et al. (2014) describe a distributed path planning algorithm that allows the robots to construct a configuration space

of the environment in a distributed fashion. A shortest-path tree is constructed using a variation of the Bellman-Ford algorithm (see Bang-Jensen and Gutin, 2008). The algorithm can cope with dynamic obstacles and changes in robot population. The algorithm is successfully tested in simulation and also with a homogeneous group of physical robots (see Habibi et al., 2015) pre-attached to an irregular object. This approach requires some robots to perform the transport while others to map the environment in order to guide the transporting robots in the direction of the goal while avoiding obstacles. While this approach is effective in selecting optimal transport trajectories, it requires the majority of robots to map the environment rather than performing the actual transport task. Yufka and Ozkan (2015) illustrate another motion planning algorithm for a group of homogeneous robots required to transport a heavy object to its final destination. This algorithm requires the robots to know their position in the environment, and also assumes that the robots can directly communicate with each other. Initially, the object trajectory is generated, and then each robot generates its trajectory to satisfy the current formation constraints. The algorithm is successfully tested with groups made of a different number of Pioneer robots pre-attached to the object.

A series of studies published in (Tuci et al., 2006; Groß and Dorigo, 2008; Gross and Dorigo, 2009) looked at the design of neuro-controllers synthesised using evolutionary computation techniques to control homogeneous groups of robots that are not required to be pre-attached to the object to be transported. The robots can physically connect both to each other and to the object. The task requires the robots to transport the object using a gripper mounted on a horizontal active axis that can be used to grasp and lift objects (Mondada et al., 2004). The robots can also change the relative orientation of the wheels with respect to the grasping point by rotating their upper body (i.e., the turret with the gripper) with respect to the chassis where the wheels are mounted. The results of these studies demonstrate that the combination of feedback generated by force sensors, the rotating turret mechanism for the effective alignment of pushing/pulling forces, and the possibility to have robot-robot connections generate an extremely effective solution to transport objects of different shapes and sizes towards a static and a moving target without the strong requirement of the robots being pre-attached to the object. In (Campo et al., 2006) and in (Ferrante et al., 2010), the collective transport strategy above mentioned has been exploited to develop two different algorithms for negotiating a common direction of transport by robots carrying an object toward a goal destination in an environment with and without obstacles.

Berman et al. (2011) try to mimic the behaviour of ants during group transport by looking for the individual rules that generate robust group-level responses. The authors observe a particular species of ant (i.e., *Aphaenogaster cockerelli*) in order to extract and reproduce in a simulated robotic system those rules that govern the ants individual actions during a foraging task requiring group transport. Individual rules are validated by comparing the behaviour of simulated and real ants. Other recent studies that follow a similar approach can be found in (Wilson et al., 2014; Gelblum et al., 2016; Guo et al., 2017).

In the remaining of this section, we review a series of studies in which the robots cooperatively transport an object on top of their

bodies. Although the robots do not have any means to directly grasp the object, we consider their transport strategies as a type of grasping strategy since the robots align their forces and sustain the transport without losing physical contact with the object as in almost the totality of the works in which robots use grasping devices to physically attach to the object to betransported.

Stilwell and Bay (1993) and Johnson and Bay (1995) describe a MRS designed to collectively transport a single palletised load. A group of simulated "ant-like" robots initially lift and then transport an item by carrying it on top of their bodies. The robots do not require an a priori knowledge about pallet mass, pallet inertia, number of the robots in the group, and their positions relative to the pallet centre of gravity. Coordination of action is achieved by sensing the forces applied to the object during the transport of the rigid pallet. In order to facilitate the dynamics of force distribution across the load, the study proposes a "reactive caster" approach that follows principles similar to those of the passive caster wheels when it aligns itself with the direction of travel. The followers align themselves with the leader by sensing the reaction force applied to their top surface with the force exerted by the leader. In (Bay, 1995), the reactive caster approach is successfully evaluated on physical robots. Any robot of the group can potentially play the role of the leader or of the follower. However, no mechanisms for a dynamic allocation of the role are contemplated. Similar studies in which a pre-selected leader manages the motion trajectory of the object and the coordination of actions is achieved through the physical interaction with the object are presented in (Pereira et al., 2002; Loh and Traechtler, 2012).

Hichri et al. (2016) propose a control algorithm in which an external server globally communicates with the robots to perform the transport task. In this study, a homogeneous group of robots is equipped with manipulators for grasping and lifting an object in order to place it on top of their bodies. Optimal positions for the robots to ensure the stability of the object are calculated based on a a priori knowledge of the number of robots in the group, object's shape, mass and centre of gravity. The server communicates position information to the robots to approach the object, to lift it, and to carry it to a destination. While carrying the object, the robots keep the desired position relative to the object, thanks to the global knowledge of the environment provided by the external server guiding the robots during transportation. The proposed approach is validated in simulation. The results of this study point to the ability of robots to maintain the stability of the object during lifting and carrying tasks. A similar approach based on the use of an external server to coordinate the robot actions is described in (Wang et al., 1994; Yamashita et al., 2003).

In (Stroupe et al., 2005) the strategy of carrying objects on robots' bodies is used in combination with a leader-follower approach. The study demonstrates, using physical robots, the capability of grasping, lifting, transporting and positioning objects in a construction task. A group of two rovers are required to manipulate objects in order to build a simple structure in a lunar-like environment. The robots communicate with each other to synchronise the grasping, lifting and placing of the objects in building the structure. The robots coordinate their actions by feeling the forces applied on the carried object using a force-torque sensor located on their manipulators. The followers coordinate with the leader by adjusting their velocity based on force-torque feedback such that the torques and forces on the manipulator remain within the experimentally defined threshold. The results indicate that the team successfully completes the construction task with a low failure rate. Another similar leader-follower cooperative transport study using a direct instead of an indirect form of communication can be found in (Hashimoto et al., 1993).

THE CAGING STRATEGY

Cooperative transport by caging is a special case of the previously discussed pushing-only strategy whereby robots intentionally entrap the object to ensure the object follows the group movements. In the caging strategy, robots arrange themselves around the object in order to form a "closure" that traps the object (Rimon and Blake, 1996). The closure must be maintained during transport to ensure the object does not escape from the robots' cage. In cooperative transport based on a caging strategy, the object's shape and size are particularly important features since they bear upon the minimum number of robots required to surround the object.

The simplest form of caging strategy with a small number of robots can be found in (Wang et al., 2004b). This study describes a variable internal force control algorithm to guide a group of three omnidirectional robots required to transport a cuboid object. The robots are cube-shaped therefore they touch the object by a line segment (rather than a point). Only the leader pushes the object while the followers hold the sides of the object tightly, such that no change occurs in the relative position and orientation between the object and each follower robot. The robots coordination is achieved by simply sensing the resultant force applied to the object and its movement. This form of indirect communication through the object is sufficient to allow the followers to maintain the formation and to contribute to the transport by exerting forces to move the object along the trajectory known only to the leader. The main limitation of this study is that the system can not follow an arbitrary trajectory that incorporates sharp turns especially when the velocity is low. Similar examples of the use of a caging strategy with the leader-follower approach can be found in (Wang et al., 2003; Wang et al., 2004a).

Brown and Jennings (1995) propose a *pusher-steerer* approach to cooperative transport which is similar to the one discussed in (Wang et al., 2004b) but without the requirement of maintaining tight contact with the object during transport. The steerer is programmed to follow a predefined path while the pusher exerts the necessary forces to transport the object. The object is placed between the pusher and the steerer. During the transport, the steerer senses the arc length travelled and adjusts its heading to follow the programmed trajectory. Using indirect communication, the pusher follows the change in the object's configuration by maintaining a fixed orientation relative to the rear face of the object. This approach is similar to a rear-wheel-drive vehicle but implemented with two separate pieces (i.e., the pusher and the steerer). The approach is validated through experiments using two physical robots to transport boxes of varying size and mass along different paths. The results indicate that the robots can successfully maintain the caging while following the programmed trajectory.

Spletzer et al. (2001) describe a cooperative transport task in which caging strategy is achieved using vision to estimate distances and relative orientations of the robots. In this study, a leader robot and two followers are required to transport an object to a destination known only to the leader. Followers maintain a desired distance and relative bearing to the leader in order to form a closure that cages the object. This approach is similar to the "pusher-watcher" approach described in (Gerkey and Matarić, 2002) and reviewed in section 2. Contrary to (Gerkey and Matarić, 2002), in this study, the watcher robot contributes to the transport by caging the object with the pushers. The main drawback of this method is that all robots need to maintain visual contact with each other.

Pereira et al. (2004) propose an algorithm for collective transport using a caging strategy that relies only on the robots' ability to estimate the object's orientation and the positions of their neighbours. In this study, three holonomic car-like robots are required to move in the direction of a goal while maintaining a formation trapping a triangular object. Each robot is equipped with an omnidirectional camera to estimate the object orientation and its position with respect to its neighbouring robots. This information is explicitly communicated between the robots to complement their partial knowledge about the object orientation and the position of other robots. The control algorithm assumes that each robot has an imaginary copy of the object attached to it at the object's origin (i.e., one of the object's corner). The intersection of these imaginary objects forms a region referred to as the *closure configuration space*. If the origin of the actual object falls inside the closure configuration space, then an object's closure is accomplished otherwise, the robots have to adjust their positions to satisfy this condition.

In a more complex scenario, (Fink et al., 2007) propose a caging strategy for a group of robots required to transport an L-shaped object on a predefined trajectory. In this study, the robots locally estimate the object closure based on direct communication regarding their position with respect to the object. Controlled by a subsumption architecture, the robots switch from the approach behaviour to the surround behaviour when they are close to the object. In surround behaviour, the robots distribute themselves around the object in order to form the potential caging. This approach requires the robot to know the object's minimum diameter (i.e., the smallest gap through which the object can fit), maximum diameter (i.e., the maximum distance between any two points in the object), and the radius of the caging circle. The robots communicate their states to neighbours until a quorum is reached—that is, when enough robots surround the object, and all are ready to initiate the transport behaviour. During transport behaviour, every robot adjusts its speed depending on the positions of neighbours and the desired trajectory of the object. If for any reason the closure is lost during transport, every robot returns to the surrounding behaviour to resume the transport. The study verifies the effectiveness of the approach using eight differential drive robots equipped with wheel encoders and laser range finders. The results of this study show the stability of the proposed caging strategy in a scenario in which the robots successfully form closures that surround the object while pushing it from an initial position to its final destination. Later in (Fink et al., 2008), the authors extend this approach to allow the robots to operate in a more complex

environment that incorporates obstacles. Another study involving a similar caging strategy generated by a subsumption architecture is described in (Eoh et al., 2014).

In (Dai et al., 2016), a control architecture based a fuzzy control methods integrated with the sliding mode method is used to control a heterogeneous group of three physical robots required to collectively transport, using a caging strategy, a convex polygon along different predefined trajectories. The robots have some predefined knowledge about the object shape. Moreover, they use a form of direct communication to share important perceptual details that help them to complement their partial knowledge of the object shape. A leader robot manages the transport by compute the inter-robot distance and required bearing of each follower. The results show that the control architecture allow the group to transport an object along different predefined linear and curved trajectories known to the leader.

Finally, the main contribution of the work described in (Wan et al., 2017) is to test a caging strategy for transporting a triangular prism to a final destination by crossing a slope terrain. In this work, the control algorithm running on a master computer generates the minimum number of robots required to securely cage the object, the initial positions of each robot with respect to the object, and the robots motion during transport. The control method exploit a direct form of communication between the master computer and the robots, and makes use of a detailed knowledge of position and orientation of the robots and the object to be transported. The study shows that simulated robots can successfully transport objects of different shape and size along the slope terrain.

DISCUSSION

The research area targeting cooperative transport by MRSs is represented by an articulated and heterogeneous body of literature. We have chosen to illustrate this literature using a categorisation system that distinguishes the research works on the basis of the type of strategy used to collectively transport the object. In this section, we illustrate general patterns that emerge from the considerable methodological diversity illustrated in previous sections, and we identify open challenges and promising directions for further work.

Our review shows that, regardless of the type of transport strategy used, a certain amount of functional diversity among the members of a group seems to be an ineluctable methodological feature to allow the robotic systems to operate in an environment with obstacles, or to develop transport trajectories that adapt to varying environmental conditions. A robot leader, or a robot watcher as in (Gerkey and Matarić, 2002), is generally deputed to direct the transport by coordinating the actions and contributions of the followers. This is a particularly recurrent pattern in those research works based on the use of grasping strategies, where the fact of having all robots attached to the object facilitates the indirect communication and allocation of duties by the leader to the followers through force sensing mechanisms. In various studies exploiting the leader-follower approach, cooperative transport is exploited to cope with objects that due to their size can be hardly transported by a single robot. However, the transport strategies generated by these heterogeneous groups tend to be very fragile

with respect to the object mass. This is because the object has to be light enough to respond to the forces exerted by the leader, that is often the only robot deputed to initiate the transport. The robustness with respect to the object mass tends to be more easily achieved by transport strategies developed by groups in which multiple agents can contribute to initiate and to sustain the transport. Another feature that tends to improve the robustness of the collective transport strategies with respect to the object mass, is the possibility for the robots to push each other as in (Fujisawa et al., 2013; Alkilabi et al., 2017), or to push and pull each other as in (Gross and Dorigo, 2009). In most of the studies we have reviewed in previous sections the robot-to-robot interactions are excluded by the use of control mechanisms designed to avoid robot-to-robot collisions. However, the exploitation of both the robot-to-robot and the robot-to-object interactions can facilitate the initial coordination of actions. Moreover, robot-to-robot interactions are particularly helpful in case of transport of heavy and relatively small objects, where the limited object perimeter prevents the group from developing the robot-to-object interactions required to initiate and sustain the transport.

When the heterogeneity of the system is the distinctive methodological feature that makes the group capable of operating in complex environments (e.g., environments with obstacles), it would be desirable if the allocation of roles or the emergence of any hierarchical organisation could be directly handled by the robots in an autonomous way. That would make possible the re-allocation of roles or the re-organisation of the group structure in case of failure of the *key element*. To the best of our knowledge, apart from the work described in (Gerkey and Matarić, 2002), heterogeneity is either based on structural differences among the robots, or on a type of functional differentiation a priori managed by the system's designer. This means that the fault-tolerance is often partially or totally sacrificed in order to boost the competencies of the group. For the future, it would be interesting to see more research works focusing on the challenge of designing mechanisms to allow a group of robots to handle functional heterogeneity by allocating, and if necessary re-allocating, important functions in a completely autonomous way. That would reconcile fault-tolerance and group competencies to design MRSs capable of carrying out complex cooperative transport tasks.

In section 3, we have intentionally included in the category of grasping strategies those research works in which MRSs transport objects on top of their bodies, even if the robots do not use any special device to attach to the object. The logic behind this choice is that for both strategies (i.e., grasping and carrying the object on top of the body) the robots align their forces and sustain the transport without losing physical contact with the object. While for robots that carry the object on top of their bodies, the persistence of the physical contact with the object during the entire transport is generally an unavoidable consequence of the way in which the robots are meant to operate (these robots generally lack any grasping device), for robots that can grasp the object with a dedicated grasping device, the physical connection could in principle be released in particular during the initial phases of the transport to facilitate the alignment of pushing/pulling forces. To the best of our knowledge, the large majority of research works where MRSs use grasping strategies to collectively transport objects

concern robots that are pre-attached to the object, and that never release the grasp during the transport (Tuci et al., 2006; Groß and Dorigo, 2008; Gross and Dorigo, 2009). The pre-attachment condition certainly takes out a large amount of complexity from research studies that tend to focus on the coordination of actions during the transport rather than on the distribution and alignment of pushing/pulling forces to initiate it. However, for the future, the development of mechanisms to allow robots to exploit both the grasp and the release action would be important not only to automate the distribution and the alignment of pushing/pulling forces, but also to improve the robustness of the system to be able to transport object of different shapes. We have seen that, apart from few studies, the large majority of the research works focus on the collective transport of rectangular objects (see **Tables 1–3**). The robustness of the collective transport strategies with respect to the object shape has been so far a rather neglected subject, that could be further investigated by researching and improving those aspects that, like the release and grasping process, directly affects it.

Methodological alternatives to develop effective transport strategies for homogeneous groups required to operate in complex environments are generally limited to solutions that work only if the object to be transported does not occlude the robots' view of the goal destination, or of the perception of eventual obstacles. The occlusion-based approach reviewed in section 2 (see Jianing Chen et al., 2015) discusses an algorithm that definitely overcomes the above-mentioned limitations and provides a very effective solution to allow homogeneous groups to cooperatively transport objects in an environment with obstacles. However, we point to the fact that, to the best of our knowledge, in the large majority of the reviewed studies, the initial alignment and the following coordination of actions is subject to the perception or to the occlusion (by the object to be transported) of the final destination of the transport (see Jianing Chen et al., 2015). This important assumption tends to simplify the initial process of the alignment of the transport forces, and largely undermines the robustness of the resulting group transport strategies to environment in which this assumption does not hold. We believe that the above mentioned assumption should be dropped to favour the robustness of the cooperative transport strategies. Finally, it is worth to note that no research work has been dedicated to the development of MRSs that can dynamically adjust the type of transport strategy (e.g., pushing, grasping, or caging) with respect to the characteristics of the object to be transported and/or of the environment in which the collective transport takes place. This is also a very interesting subject for future work.

CONCLUSIONS

We have reviewed the literature on MRSs focused on the development of hardware and control systems to allow autonomous robots to cooperatively transport objects that can not be moved by a single robot. We have structured our review on a rather unconventional and relatively "coarse-grained" categorisation framework based on the type of transport strategy used by the robotic systems to move the objects. With this framework, we have ordered a rather heterogeneous body of MRSs literature, by focusing not only on motivations and objectives, but also on those distinctive

TABLE 3 | Main characteristics of research works using a caging transport strategy.

	Experimental setting						Control and communication		Similar studies
	nr. robots	group type: -homogeneous (ho) -heterogeneous (he)	autonomous allocation of roles	env. with obstacles	obj. shape	transport direction	controller design method	communication	
Wang et al., 2004b	3	he	no	no	rectangular prism	known to leader	other	indirect	Wang et al., 2003; Wang et al., 2004a
Brown and Jennings, 1995	2	he	no	no	rectangular prism	known to leader	other	indirect	
Spletzer et al., 2001	3	he	no	no	rectangular prism	known to leader	other	indirect	
Pereira et al., 2004	3–4	ho	no	no	triangular prism	towards a goal known to all robots	other	direct	Wan et al., 2012
Fink et al., 2007	4–8	ho	no	yes	L-shaped	towards a goal known to all robots	behaviour-based	indirect	Fink et al., 2008; Eoh et al., 2014
Dai et al., 2016	3	he	no	no	convex polygon	predefined trajectories known to the leader	other	direct	
Wan et al., 2017	3	ho	no	no	triangular prism	predefined trajectories	other	direct	

See also caption Table 1 for further illustration of the descriptive categories.

methodological details that characterise the contribution of each single reviewed study. In section 5, we have critically examined common features emerging from the comparison of within and between categories works, and we have pointed to potentially fruitful directions for future work.

We wish to conclude this review with a brief reference to cooperative transport in natural systems. Ants have evolved extremely effective competencies to cooperatively retrieve items that can be hundreds or even thousands times the weight an individual can carry (Czaczkes et al., 2011). Owing to cooperative transport, ants can perform faster prey retrieval reducing both the exposition of foragers to predators, and the risk of food being caught and eaten by other aggressive species (Hölldobler et al., 1978; Yamamoto et al., 2009). The speedy retrieval of prey also reduces the time workers are involved in transport tasks, freeing them for other colony relevant tasks (Feener and Moss, 1990; Tanner, 2008). Cooperative transport also reduces the energy cost of transport by allowing carriers to keep up with the dense flow of traffic and by reducing the possibility of traffic jams (Czaczkes and Ratnieks, 2013). Biologists suggest that these complex group level responses are underpinned by simple behavioural rules (Franks, 1986; McCreery et al., 2016). We think that important lessons can still be learned from observing the complex cooperative transport behaviour shown by various ant species. It is then the task of roboticists to transform these observations into fruitful design principles and effective methodological choices to develop robust, flexible and scalable MRSs that cooperatively transport objects.

AUTHOR CONTRIBUTIONS

The authors have equally contributed to the analysis of the literature and to the writing of the paper.

FUNDING

REFERENCES

Alkilabi, M., Narayan, A., Lu, C., and Tuci, E. (2016). "Evolving group transport strategies for e-puck robots: moving objects towards a target area," in *Proc. of the Int. Symposium on Distributed Autonomous Robotic Systems (DARS)*, ed. R. G (Germany: Springer).

Alkilabi, M. H. M., Narayan, A., and Tuci, E. (2017). Cooperative object transport with a swarm of e-puck robots: robustness and scalability of evolved collective strategies. *Swarm Intell.* 11 (3-4), 185–209. doi: 10.1007/s11721-017-0135-8

Bahçeci, E., Soysal, O., and Şahin, E. (2003). *A review: pattern formation and adaptation in multi-robot systems*. Pittsburgh, PA: Robotics Institute, Carnegie Mellon University.

Bang-Jensen, J., and Gutin, G. (2008). *Digraphs: Theory, Algorithms and Applications*, second Edn. Germany: Springer.

Bay, J. S. (1995). Design of the "army-ant" cooperative lifting robot. *IEEE Robot. Automat. Mag.* 2 (1), 36–43. doi: 10.1109/100.388293

Bayindir, L., and Şahin, E. (2007). A review of studies in swarm robotics. *Turkish Journal of Electrical Engineering & Computer Sciences* 15, 115–147.

Berman, S., Lindsey, Q., Sakar, M. S., Kumar, V., and Pratt, S. C. (2011). Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems. *Proc. IEEE* 99 (9), 1470–1481. doi: 10.1109/JPROC.2011.2111450

Bernard, M., Kondak, K., Maza, I., and Ollero, A. (2011). Autonomous transportation and deployment with aerial robots for search and rescue missions. *J. Field Robotics* 28 (6), 914–931. doi: 10.1002/rob.20401

Bicho, E., and Schöner, G. (1997). The dynamic approach to autonomous robotics demonstrated on a low-level vehicle platform. *Rob. Auton. Syst.* 21 (1), 23–35. doi: 10.1016/S0921-8890(97)00004-3

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.* 7 (1), 1–41. doi: 10.1007/s11721-012-0075-2

Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE J. Robot. Automat.* 2 (1), 14–23. doi: 10.1109/JRA.1986.1087032

Brown, R., and Jennings, J. (1995). A pusher/steerer model for strongly cooperative mobile robot manipulation. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots* 3, 562–568.

Cai, Y., and Yang, S. (2012). "A survey on multi-robot systems," in *IEEE World Automation Congress (WAC)*, Vol. 2012 1–6.

Campo, A., Nouyan, S., Birattari, M., Groß, R., and Dorigo, M. (2006). "Negotiation of goal direction for cooperative transport," in *International Workshop on Ant Colony Optimization and Swarm Intelligence*, eds M. Dorigo, L. Gambardella, M. Birattari, A. Martinoli, R. Poli, and T. Stützle 191–202.

Cao, Y. U., Fukunaga, A. S., and Kahng, A. (1997). Cooperative mobile robotics: antecedents and directions. *Auton. Robots* 4 (1), 7–27. doi: 10.1023/A:1008855018923

Czaczkes, T. J., Nouvellet, P., and Ratnieks, F. L. W. (2011). Cooperative food transport in the Neotropical ant, *Pheidole oxyops*. *Insectes Soc.* 58 (2), 153–161. doi: 10.1007/s00040-010-0130-1

Czaczkes, T., and Ratnieks, F. (2013). Cooperative transport in ants (hymenoptera: Formicidae) and elsewhere. *Myrmecological News* 18, 1–11.

Dai, Y., Kim, Y., Wee, S., Lee, D., Lee, S., and S.Lee. (2016). Symmetric caging formation for convex polygonal object transportation by multiple mobile robots based on fuzzy sliding mode control. *ISA Trans.* 60, 321–332. doi: 10.1016/j.isatra.2015.11.017

Emery, R., Sikorski, K., and Balch, T. (2002). "Protocols for collaboration, coordination and dynamic role assignment in a robot team," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Vol. 3 3008–3015.

Eoh, G., Lee, S. -H., Lee, T. -S., and Lee, B. H. (2014). Distributed object transportation using virtual object. *JIII* 2 (1), 20–25. doi: 10.12720/jiii.2.1.20-25

Farinelli, A., Iocchi, L., and Nardi, D. (2004). Multirobot systems: a classification focused on coordination. *IEEE Trans. Syst. Man Cybern. B Cybern.* 34 (5), 2015–2028. doi: 10.1109/TSMCB.2004.832155

Farivarnejad, H., Wilson, S., and Berman, S. (2016). "Decentralized sliding mode control for autonomous collective transport by multi-robot systems," in *Proc. of the IEEE Decision and Control Conference (CDC)* 1826–1833.

Feener, D., and Moss, K. G. (1990). Defense against parasites by hitchhikers in leaf-cutting ants: a quantitative assessment. *Behav. Ecol. Sociobiol. (Print).* 26 (1), 17–29. doi: 10.1007/BF00174021

Ferrante, E., Brambilla, M., Birattari, M., and Dorigo, M. (et al2010). "Socially-mediated negotiation for obstacle avoidance in collective transport," in *Proc. of Int. Symposium on Distributed Autonomous Robotic Systems*, Vol. 83, eds A. Martinoli, F. Mondada, N. Correll, G. Mermoud, M. Egerstedt, M. Hsieh, L. Parker, K. Stöy., et al. (Germany: Springer), 571–583.

Fink, J., Hsieh, M., and Kumar, V. (2008). "Multi-robot manipulation via caging in environments with obstacles," in *Proc. of the IEEE Int. Conf. on Robotics and Automation(ICRA)* 1471–1476.

Fink, J., Michael, N., and Kumar, V. (2007). "Composition of vector fields for multi-robot manipulation via caging," in *In Proc. of the Conf. Robotics: Science and Systems*, Vol. 3.

Franks, N. R. (1986). Teams in social insects: group retrieval of prey by army ants (Eciton burchelli, Hymenoptera: Formicidae). *Behav. Ecol. Sociobiol. (Print).* 18 (6), 425–429. doi: 10.1007/BF00300517

Fujisawa, R., Imamura, H., and Matsuno, F. (2013). "Cooperative transportation by swarm robots using pheromone communication," in *Proc. of the INt.*

Symposium on Distributed Autonomous Robotic Systems (DARS). Germany: Springer, 559–570.

Gelblum, A., Pinkoviezky, I., Fonio, E., Gov, N. S., and Feinerman, O. (2016). Emergent oscillations assist obstacle negotiation during ant cooperative transport. *Proc. Natl. Acad. Sci. U.S.A.* 113 (51), 14615–14620. doi: 10.1073/pnas.1611509113

Gerkey, B., and Matarić, M. (2002). "Pusher-watcher: An approach to fault-tolerant tightly-coupled robot coordination," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 1 464–469.

Gerkey, B., and Matarić, M. (2001). "Principled communication for dynamic multi-robot task allocation," in *Experimental Robotics VII*, eds D. Rus, and S. Sanjiv (Germany: Springer), 353–362.

Gross, R., and Dorigo, M. (2009). Towards group transport by swarms of robots. *IJBIC* 1 (1/2), 1–13. doi: 10.1504/IJBIC.2009.022770

Groß, R., and Dorigo, M. (2008). Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling. *Adapt. Behav.* 16, 285–305.

Groß, R., and Dorigo, M. (2004). "Cooperative transport of objects of different shapes and sizes," in *Proc. of the Int. Workshop on Ant Colony Optimization and Swarm Intelligence*, Vol. 3172, eds M. Dorigo, M. Birattari, C. Blum, L. M. Gambardella, F. Mondada, and T. Stützle (Germany: Springer), 106–117.

Guo, S., Wang, M., Notman, G., Chang, J., Zhang, J., and Liao, M. (2017). Simulating collective transport of virtual ants. *Comput. Animat. Virtual Worlds* 28 (3-4):e1779. doi: 10.1002/cav.1779

Habibi, G., Kingston, Z., Xie, W., Jellins, M., and McLurkin, J. (2015). "Distributed centroid estimation and motion controllers for collective transport by multi-robot systems," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, ed. A. O 1282–1288.

Habibi, G., Xie, W., Jellins, M., and McLurkin, J. (2014). "Distributed path planning for collective transport using homogeneous multi-robot systems," in *Proc. of the Int. Symposium on Distributed Autonomous Robotic Systems (DARS)*, eds M. A. Hsieh, and G. Chirikjian (Germany: Springer), 151–164.

Hashimoto, M., Oba, F., and Eguchi, T. (1993). "Dynamic control approach for motion coordination of multiple wheeled mobile robots transporting a single object," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Vol. 3 1944–1951.

Hekmatfar, T., Masehian, E., and Mousavi, S. J. (2014). "Cooperative object transportation by multiple mobile manipulators through a hierarchical planning architecture," in *Second RSI/ISM Int. Conf. on Robotics and Mechatronics (ICRoM*, ed. editor (), 503–508IEEE.

Hichri, B., Adouane, L., Fauroux, J., Mezouar, Y., and Doroftei, I. (2016). "Cooperative mobile robot control architecture for lifting and transportation of any shape payload," in *Proc. of the Int. Symposium on Distributed Autonomous Robotic Systems (DARS)* 177–191.

Hu, W., Ishii, K. S., and Ohta, A. T. (2011). Micro-assembly using optically controlled bubble microrobots. *Appl. Phys. Lett.* 99 (9):094103. doi: 10.1063/1.3631662

Huntsberger, T., Rodriguez, G., and Schenker, P. (2000). "Robotics challenges for robotic and human mars exploration," in *Proc. of Robotics 2000*, ed. W. Stone (United States: ASCE), 340–346.

Hölldobler, B., Stanton, R. C., and Markl, H. (1978). Recruitment and food-retrieving behavior in Novomessor (Formicidae, Hymenoptera). *Behav. Ecol. Sociobiol. (Print).* 4 (2), 163–181. doi: 10.1007/BF00354978

Iocchi, L., Nardi, D., and Salerno, M. (2000). "Reactivity and deliberation: a survey on multi-robot systems" *Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems* 9–32.

Jiang, S., Cao, J., Liu, Y., Chen, J., and Liu, X. (2016). "Programming large-scale multi-robot system with timing constraints," in *Proc. of the Int. Conf. on Computer Communication and Networks (ICCCN)* 1–9.

Jianing Chen ., Gauci, M., Wei Li ., Kolling, A., Gros, R, Chen, J . (2015). Occlusion-based cooperative transport with a swarm of miniature mobile robots. *IEEE Trans. Robot.* 31 (2), 307–321. doi: 10.1109/TRO.2015.2400731

Johnson, P. J., and Bay, J. S. (1995). Distributed control of simulated autonomous mobile robot collectives in payload transportation. *Auton. Robots* 2 (1), 43–63. doi: 10.1007/BF00735438

Kapellmann-Zafra, G., Chen, J., and Groß, R. (2016). "Using google glass in human–robot swarm interaction," in *Proc. of the Conf. Towards Autonomous Robotic Systems (TAROS)* 196–201.

Khamis, A., Hussein, A., and Elmogy, A. (2015). "Multi-robot task allocation: A review of the state-of-the-art," in *Cooperative Robots and Sensor Networks 2015*, eds A. Koubâa, and J. Martínez-de Dios (Germany: Springer), 31–51.

Kosuge, K., and Oosumi, T. (1996). "Decentralized control of multiple robots handling an object," in *Proc. of the IEEE/RSJ Int. Confe. on Intelligent Robots and Systems (IROS)*, Vol. 1 318–323.

Kosuge, K., Oosumi, T., Satou, M., Chiba, K., and Takeo, K. (1998). "Transportation of a single object by two decentralized-controlled nonholonomic mobile robots," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Vol. 4 2989–2994.

Kube, C. R., and Bonabeau, E. (2000). Cooperative transport by ants and robots. *Rob. Auton. Syst.* 30 (1-2), 85–101. doi: 10.1016/S0921-8890(99)00066-4

Kube, C. R., Hong Zhang , ., and . (1993). Collective Robotics: From Social Insects to Robots. *Adapt. Behav.* 2 (2), 189–218. doi: 10.1177/105971239300200204

Kube, C., and Zhang, H. (1997). "Task modelling in collective robotics," in *Robot Colonies*, eds R. Arkin, and G. Bekey (USA: Springer), 53–72.

Loh, C. C., and Traechtler, A. (2012). Cooperative transportation of aload using nonholonomic mobile robots. *Procedia Engineering* 41, 860–866. doi: 10.1016/j.proeng.2012.07.255

Machado, T., Malheiro, T., Monteiro, S., Erlhagen, W., and Bicho, E. (2016). "Multi-constrained joint transportation tasks by teams of autonomous mobile robots using a dynamical systems approach" *Proceedings of the IEEE Int. Conf. on Robotics and Automation (ICRA)* 3111–3117.

Martín H., J. A., de Lope, J., and Maravall, D. (2010). Analysis and solution of a predator–protector–prey multi-robot system by a high-level reinforcement learning architecture and the adaptive systems theory. *Rob. Auton. Syst.* 58 (12), 1266–1272. doi: 10.1016/j.robot.2010.08.005

Mataric, M., Nilsson, M., and Simsarin, K. (1995). "Cooperative multi-robot box-pushing," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Vol. 3 556–561.

McCreery, H. F., Correll, N., Breed, M. D., and Flaxman, S. (2016). Consensus or deadlock? Consequences of simple behavioral rules for coordination in group decisions. *PLoS ONE* 11 (9):e0162768. doi: 10.1371/journal.pone.0162768

Michael, N., Fink, J., and Kumar, V. (2011). Cooperative manipulation and transportation with aerial robots. *Auton. Robots* 30 (1), 73–86. doi: 10.1007/s10514-010-9205-0

Mondada, F. (2009). "The e-puck, a robot designed for education in engineering" *Proc. of the Int. Conf. on Autonomous Robot Systems and Competitions (IEEE)* 59–65.

Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. W., Floreano, D., Deneubourg, J.-L., et al. (2004). Swarm-bot: a new distributed robotic concept. *Auton. Robots* 17 (2/3), 193–221. doi: 10.1023/B:AURO.0000033972.50769.1c

Neumann, M., Chin, M., and Kitts, C. (2014). "Object manipulation through explicit force control using cooperative mobile multi-robot systems" *Proceedings of the World Congress on Engineering and Computer Science* 364–369.

Parker, L. E. (1998). Alliance: an architecture for fault tolerant multirobot cooperation. *IEEE Trans. Robot. Automat.* 14 (2), 220–240. doi: 10.1109/70.681242

Parker, L. (2000). Lifelong adaptation in heterogeneous multi-robot teams: response to continual variation in individual robot performance. *Auton. Robots* 2000, 239–267.

Parker, C. A. C., and Zhang, H. (2006). Collective Robotic Site Preparation. *Adapt. Behav.* 14 (1), 5–19. doi: 10.1177/105971230601400101

Pereira, G. A. S., Campos, M. F. M., and Kumar, V. (2004). Decentralized algorithms for multi-robot manipulation via caging. *Int. J. Rob. Res.* 23 (7-8), 783–795. doi: 10.1177/0278364904045477

Pereira, G., Pimentel, B., Chaimowicz, L., and Campos, M. (2002). "Coordination of multiple mobile robots in an object carrying task using implicit communication," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Vol. 1 281–286.

Quinn, M. (2004) The evolutionary design of controllers for minimally-equipped homogeneous multi-robot systems. Ph.D. thesis. Brighton: University of Sussex.

Rahman, M. A., Cheng, J., Wang, Z., and Ohta, A. T. (2017). Cooperative Micromanipulation Using the Independent Actuation of Fifty Microrobots in Parallel. *Sci. Rep.* 7 (1):3278. doi: 10.1038/s41598-017-03525-y

Rimon, E., and Blake, A. (1996). Caging 2d bodies by 1 parameter two fingered gripping systems. In *Proc. of the IEEE Int. Conf. on Robotics and Automation* 2, 1458–1464.

Şahin, E. (2005). "Swarm robotics: From sources of inspiration to domains of application," in *Proc. of the First Int. Workshop on Swarm Robotics*, eds E. Şahin, and W. Spears (Springer) 10–20.

Sasaki, J., Ota, J., Yoshida, E., Kurabayashi, D., and Arai, T. (1995). "Cooperating grasping of a large object by multiple mobile robots," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (TCRA)*, Vol. 1 1205–1210.

Sen, S., Sekaran, M., and Hale, J. (1994). "Learning to coordinate without sharing information" *Proc. of the National Conf. on Artificial Intelligence, AAAI-94* 426–431.

Shahrokhi, S., and Becker, A. (2016). "title," in *IEEE International Conference on Automation Science and Engineering (CASE)*, ed. editor (IEEE) 561–566.

Simmons, R., Singh, S., Hershberger, D., Ramos, J., and Smith, T. (2001). "First results in the coordination of heterogeneous robots for large-scale assembly," in *Experimental Robotics VII*, eds D. Rus, and S. Sanjiv (Germany: Springer), 323–332.

Soares, R., Bicho, E., Machado, T., and Erlhagen, W. (2007). "Object transportation by multiple mobile robots controlled by attractor dynamics: Theory and implementation" *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)* 937–944.

Spletzer, J., Das, A., Fierro, R., Taylor, C., Kumar, V., and Ostrowski, J. (2001). "Cooperative localization and control for multi-robot manipulation," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Vol. 2 631–636.

Stilwell, D., and Bay, J. (1993). "Toward the development of a material transport system using swarms of ant-like robots," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)* 766–771.

Stroupe, A., Huntsberger, T., Okon, A., Aghazarian, H., and Robinson, M. (2005). "Behavior-based multi-robot collaboration for autonomous construction tasks," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)* 1495–1500.

Sugie, H., Inagaki, Y., Aisu, SOH., and Unemi, T. (1995). "Placing objects with multiple mobile robots-mutual help using intention inference" *Proc. of the IEEE Int. Conf. on Robotics and Automation* 2181–2186.

Takeda, H., Hirata, Y., Wang, Z., and Kosuge, K. (2002). "Collision avoidance algorithm for two tracked mobile robots transporting a single object in coordination based on function allocation concept," in *Distributed Autonomous Robotic Systems 5*, eds H. Asama, T. Arai, T. Fukuda, and T. Hasegawa (Germany: Springer), 155–164.

Tang, F., and Parker, L. (2005). "Asymtre: Automated synthesis of multi-robot task solutions through software reconfiguration," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)* 1501–1508.

Tanner, C. J. (2008). Resource characteristics and competition affect colony and individual foraging strategies of the wood ant Formica integroides. *Ecol. Entomol.* 33 (1), 127–136. doi: 10.1111/j.1365-2311.2007.00950.x

Tuci, E., Gross, R., Trianni, V., Mondada, F., Bonani, M., and Dorigo, M. (2006). Cooperation through self-assembly in multi-robot systems. *ACM Trans. Auton. Adapt. Syst.* 1 (2), 115–150. doi: 10.1145/1186778.1186779

Wan, W., Rui, F., Masamichi, S., Tomomasa, S., and Yasuo, K. (2012). "Cooperative manipulation with least number of robots via robust caging" *IEEE/ASME Int. Conf. on Advanced Intelligent Mechatronics (AIM)* 896–903.

Wan, W., Shi, B., Z.Wang., and Fukui, R. (2017). "Multirobot object transport via robust caging" *IEEE Transactions on Systems, Man, and Cybernetics: Systems*

Wang, Y., and de Silva, C. (2006a). "Cooperative transportation by multiple robots with machine learning" *Proc. of the IEEE Congress on Evolutionary Computation (CEC)* 3050–3056.

Wang, Y., and de Silva, C. (2006b). "Multi-robot box-pushing: Single-agent q-learning vs. team q-learning" *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)* 3694–3699.

Wang, Z., Hirata, Y., and Kosuge, K. (2003). "Control multiple mobile robots for object caging and manipulation," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Vol. 2 1751–1756.

Wang, Z., Hirata, Y., and Kosuge, K. (2004a). "Control a rigid caging formation for cooperative object transportation by multiple mobile robots," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Vol. 2 1580–1585.

Wang, Z., Nakano, E., and Matsukawa, T. (1994). "Cooperating multiple behavior-based robots for object manipulation," in *Proc. of the IEEE/RSJ/GI Int. Conf. on Intelligent Robots and Systems*, Vol. 3 1524–1531.

Wang, Z., and Schwager, M. (2016). "Kinematic multi-robot manipulation with no communication using force feedback" *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)* 427–432.

Wang, Z., Takano, Y., Hirata, Y., and Kosuge, K. (2004b). "A pushing leader based decentralized control method for cooperative object transportation," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Vol. 1 1035–1040.

Wang, Z., Yang, G., Su, X., and Schwager, M. (2016). "Ouijabots: Omnidirectional robots for cooperative object transport with rotation control using no communication," in *Proc. on the Int. Symposium of Distributed Autonomous Robotic Systems (DARS)*. Germany: Springer, 1–10.

Wilson, S., Gameros, R., Sheely, M., Lin, M., Dover, K., Gevorkyan, R., et al. (2016). Pheeno, A Versatile Swarm Robotic Research and Education Platform. *IEEE Robot. Autom. Lett.* 1 (2), 884–891. doi: 10.1109/LRA.2016.2524987

Wilson, S., Pavlic, T. P., Kumar, G. P., Buffin, A., Pratt, S. C., and Berman, S. (2014). Design of ant-inspired stochastic control policies for collective transport by robotic swarms. *Swarm Intell.* 8 (4), 303–327. doi: 10.1007/s11721-014-0100-8

Woern, H., Szymanski, M., and Seyfried, J. (2006). "The i-swarm project," in *Proc. of the IEEE Int. Symposium on Robot and Human Interactive Communication*, ed. K. D (IEEE) 492–496.

Yamada, S., and Saito, J. (2001). Adaptive action selection without explicit communication for multirobot box-pushing. *IEEE Trans. Syst. Man Cybern. C* 31 (3), 398–404. doi: 10.1109/5326.971668

Yamamoto, A., Ishihara, S., Ito, F., and Fuminori, I. (2009). Fragmentation or Transportation: Mode of Large-Prey Retrieval in Arboreal and Ground Nesting Ants. *J. Insect Behav.* 22 (1), 1–11. doi: 10.1007/s10905-008-9126-3

Yamashita, A., Arai, T., Jun Ota ., and Asama, H. (2003). Motion planning of multiple mobile robots for cooperative manipulation and transportation. *IEEE Trans. Robot. Automat.* 19 (2), 223–237. doi: 10.1109/TRA.2003.809592

Yan, Z., Jouandeau, N., and Cherif, A. A. (2013). A Survey and Analysis of Multi-Robot Coordination. *Int. Journal of Advanced Robotic Systems* 10 (12), 399. doi: 10.5772/57313

Yufka, A., and Ozkan, M. (2015). Formation-Based Control Scheme for Cooperative Transportation by Multiple Mobile Robots. *International Journal of Advanced Robotic Systems* 12 (9), 120. doi: 10.5772/60972

Opinion Dynamics With Mobile Agents: Contrarian Effects by Spatial Correlations

Heiko Hamann*

Institute of Computer Engineering, University of Lübeck, Lübeck, Germany

We investigate the dynamics of opinion formation in a group of mobile agents with noisy perceptions. Two models are applied, the 2-state Galam opinion dynamics model with contrarians and an urn model of collective decision-making. It is shown that models built on the well-mixed assumption fail to represent the dynamics of a simple scenario. The challenge of accounting for correlations in the agents' spatial distribution is overcome by different heuristics and supported by empirical investigations. We present a concise, simple 1-dimensional macroscopic modeling approach that can be tuned to correctly model spatial correlations.

Keywords: swarm robotics, swarm intelligence, opinion dynamics, collective decision making, swarm robotic system

1. INTRODUCTION

Group behaviors of interacting, mobile agents are of interest in many fields and many models have been published. So-called microscopic models (also known as multi-agent models, agent-based models, or individual-based models) explicitly incorporate properties of each member of the group such as position, direction, and internal state. Examples are models of self-propelled particles (Vicsek et al., 1995; Czirók and Vicsek, 2000; Levine et al., 2000) and active Brownian agents (Schimansky-Geier et al., 1995; Helbing et al., 1997; Schweitzer, 2003). So-called macroscopic models abstract away such individual properties (e.g., derivations in the mean-field limit) and reduce the state space to a few variables. Examples are diffusion models of animal groups (Okubo, 1986; Hillen and Painter, 2009; Degond and Yang, 2010; Vicsek and Zafeiris, 2012), robots (Galstyan et al., 2005; Hamann, 2010, 2018; Prorok et al., 2011), and general models of self-propelled particles (Czirók and Vicsek, 2000). Collective decision-making, in particular, is observed in many systems such as natural swarms (Franks et al., 2003; Nicolis and Dussutour, 2008; Yates et al., 2009), artificial swarms (Schmickl et al., 2008; Garnier et al., 2009), and in human groups and societies (Galam and Moscovici, 1991; Helbing and Molnar, 1997; Hegselmann and Krause, 2002; Galam, 2004; Galam and Jacobs, 2007; Motsch and Tadmor, 2014). Naturally, observations and descriptions of these systems take place on two different levels: the microscopic level, where an individual agent is observed and described, and the macroscopic level, where the group of agents is considered as a whole. This categorization holds also for models of opinion dynamics. Microscopic models represent internal states and in the case of spatial models also positions of each agent which increases the computational effort that is to be invested to evaluate the model. In macroscopic models one abstracts from details of individual agents, for example, in a mean-field approach (Schweitzer, 2002), and tries to focus on important macroscopic features. The macroscopic models are the epistemologically more promising approach because they allow for deeper insights as stated by Schweitzer (2003):

"To gain insight into the interplay between microscopic interactions and macroscopic features, it is important to find a level of description that, on the one hand, considers specific features of the system and is suitable for reflecting the origination of new qualities, but, on the other hand, is not flooded with microscopic details."

There are macroscopic models that are built on simplifying assumptions, for example, there are models of opinion dynamics that assume well-mixed agent distributions (Schweitzer et al., 2002; Galam, 2004), that is, uniform distributions of agents independent of their current opinion. While it is possible, for example, to derive a Fokker-Planck equation of Brownian motion with drift based on integration over short time intervals assuming uncorrelated collisions of particles (Haken, 1977), it is in general not possible for biological swarm models due to the breakdown of the "propagation of chaos" (Carlen et al., 2013).

A frequently used method to incorporate spatial correlations of agents and interactions (Mateo et al., 2017), be it due to spatial relations or relations based on opinions, is that of voter models based on networks. Opinion dynamics models and swarm models have both two different types: discrete (Sood and Redner, 2005; Holme and Newman, 2006) and continuous (Toner and Tu, 1998). Whether spatiality is of importance in swarm and opinion dynamics models is questioned. For example Huepe et al. (2011) argue that

"spatial geometry may have less of an impact on collective motion than previously thought."

A simple modeling approach is based on so-called "adaptive coevolutionary networks" which are of low dimension and non-spatial (Gross and Blasius, 2008; Huepe et al., 2011).

We consider the well-mixed assumption as too imprecise for certain applications (Hamann, 2012, 2013) because agent distributions might be intrinsically correlated and consequently models based on the well-mixed assumption have limited accuracy. These applications, such as collective motion of locusts (Yates et al., 2009), hung elections (Galam, 2004), or aggregation behaviors of robot swarms (Schmickl and Hamann, 2011), are of importance. Hence, we assume that spatial correlations exist but we also want to restrict ourselves to very concise and easy to handle models of low dimensions. The motivation of this paper is to show how the limitations of the well-mixed assumption can be overcome while still keeping the models concise and easily manageable.

In the following, we investigate a binary decision problem in a group of mobile agents with noisy perceptions and compare results of two opinion dynamics models: first, the 2-state Galam opinion dynamics model with contrarians (Galam, 2004) and, second, an urn model for collective decisions in swarms (Hamann, 2012). The Galam model is particularly suited to the investigated multi-agent system because it accounts for the size of subgroups, that influence each others' opinion, which is also explicitly set in the multi-agent system. However, it does not account for spatial correlations between agents. The urn

model is of interest because it allows for a description of spatial correlations but has no concept of subgroup sizes.

The multi-agent system that is investigated here was introduced before (Hamann and Wörn, 2007; Hamann et al., 2010) and was labeled "density classification scenario" because the agents' choice is set close to a symmetric setting initially and the supposed task is that all agents should converge to the choice that had a slight majority initially. Here, we are not interested in collective decision-making as such but only the spatial correlations of the opinion dynamics. The agents show a simple form of motion. They move like billiard balls without friction. They move straight within a square and bounce off each other and the bounding walls.

2. DENSITY CLASSIFICATION SCENARIO

In this scenario, we have a population of N agents that are in one of two states: either they are in favor of opinion A or in favor of opinion B. Originally this scenario is interpreted as a task that is assigned to a population to estimate whether there are initially more A- or more B-members, that is, to converge on a majority decision (Hamann and Wörn, 2007; Hamann et al., 2010). This problem is derived from a well-known example of emergent computation in cellular automata (Packard, 1988).

We define this system as a 2-d self-propelled particles model. The particles move in a bounded square of dimensionless side length 1 (unit square). Collisions between particles and bounds are elastic.

Paricles also avoid collisions with each other by bouncing off as soon as they are within a collision avoidance radius $r = 0.01$.[1] All particles have equal velocity of 0.01 at all times (see **Table 1** for all parameters). Particle positions $\mathbf{x}(t)$ and states $o(t)$ have initially a random uniform distribution (i.e., initial positions sampled from a uniform distribution; 50% of agents in favor of A, 50% in favor of B).

We include an explicit stochastic component because we assume errors in the opinion recognition. We assume that a particle recognizes the state of an encountered particle correctly only with a given probability $1 - \gamma = 0.8$. A particle perceives the state of particle j as

$$p(o_j(t)) = \begin{cases} o_j(t), & \text{with probability } 1 - \gamma \\ \overline{o_j}(t), & \text{with probability } \gamma \end{cases}, \qquad (1)$$

whereas $\overline{o_j}$ is the opposite of the opinion of particle j.

The particles have an internal memory $\mathcal{N}$. Whenever at least two particles i and j are mutually within perception range $r = 0.01$ ($\|\mathbf{x}_j(t) - \mathbf{x}_i(t)\| \leqslant r$), they perceive the opinion of each other ($p(o_i(t))$ and $p(o_j(t))$ respectively), and store it in their memory $\mathcal{N}_i \cap \{p(o_j(t))\}$ an $\mathcal{N}_j \cap \{p(o_i(t))\}$ respectively. Once a particle had $|\mathcal{N}| = 5$ of these particle–particle encounters[2], it

---

[1]Note that the distance $r = 0.01$ is not always enforced because of the particles' high velocity of 0.01 per time step. Distances between particles below $r = 0.01$ do occur. Once such an event is detected particles turn away from each other.
[2]Choice of five is arbitrary, while odd numbers are preferred to avoid tie-breaking methods; the agent does not change its opinion until five encounters have occurred;

Parameter	Value
Swarm size N	150
Square world side length	1
Avoidance distance	0.01
Agent speed per step	0.01
Iterations per simulation run	8×10^3
Repetitions M (independent sim. runs)	1×10^6

reconsiders its current opinion, converts to the opinion that was more frequent in these five encounters, and resets its memory to $\mathcal{N} = \emptyset$. The above given parameters are set as stated in **Table 1**, such that a particle does not travel far (i.e., only fractions of the unit square) to gather five opinions. Hence, the system is not necessarily well-mixed, there is a chance for spatial correlations to form, and a particle's memory $\mathcal{N}$ can be interpreted as its perception of its neighborhood.

3. MODELING APPROACH I: GALAM MODEL

In the following we apply the 2-state Galam opinion dynamics model (Galam and Moscovici, 1991; Galam, 1997, 2008). It is a non-spatial model with discrete time and based on a population of N agents. In each round, agents come together in small groups of size m that are randomly picked without any bias. Within these groups a local majority rule is applied (i.e., the whole group switches to the group's majority opinion). If m is odd, tie breakers need not to be considered.

The density classification task is similar to the 2-state Galam opinion dynamics model concerning the decision process which is based on observing five particles and subsequently switching to the state of the majority. However, the formation of these virtual groups is neither necessarily mutual due to asynchronous decisions nor uncorrelated due to the spatial distribution of particles. Still, we apply Galam's model as an approximation. We set the group size to $m = 5$. In addition we apply Galam's extension of his model, the so-called "contrarians" (Galam, 2004). Galam's assumption is that a fraction a of the population are contrarians, that is, they always switch to the minority opinion of their group. We use the contrarian concept here to model effects due to spatial correlations, which will become clear in the following.

The model is based on one state variable s_t. Say we count a number of A_t agents with opinion A at time t, then we define the global opinion state $s_t = A_t/N$ which is the fraction of the population with opinion A. The dynamics of the state variable s_t

_____

if there are more neighbors, a random subset is chosen to keep the limit of five encounters.

according to the 2-state Galam opinion dynamics model with contrarians for a group size of $m = 5$ is

$$
\begin{aligned}
s_{t+1} = g(s_t, a) = &(1-a)(10s_t^3(1-s_t)^2 + 5s_t^4(1-s_t) + s_t^5) \\
&+ a(10(1-s_t)^3(1-(1-s_t))^2 + 5(1-s_t)^4(1-(1-s_t)) \\
&+ (1-s_t)^5).
\end{aligned} \tag{2}
$$

This model is based on simple probability theory and combinatorics, for details see Galam (2004). In **Figure 1A** we give a plot for $\Delta s_t = s_{t+1} - s_t = g(s_t, a) - s_t$ with $a = 0$. For $g(s_t, a = 0)$ we have two stable fixed points ($s_1^* = 0$ and $s_2^* = 1$)[3]. Due to the noisy perception of particles in the density classification task, this does not well correspond to the observation in the simulations. Even for $s = 0$ agents will on average still perceive an effective state of $s = \gamma$ (discussed below concerning **Figure 1B**).

Next, we want to empirically investigate the spatial correlations between particles in the density classification simulation. We define the local perception $s_i^{\mathrm{loc}}(s, t)$ of the global state s by a particle i as

$$
s_i^{\mathrm{loc}}(s, t) = \begin{cases} \frac{1}{|\mathcal{N}_i(t)|}|\{o|o = A, o \in \mathcal{N}_i(t)\}|, & |\mathcal{N}_i(t)| > 0 \\ \text{undefined}, & |\mathcal{N}_i(t)| = 0 \end{cases}. \tag{3}
$$

In order to get statistically useful measurements, we define the local opinion state $s^{\mathrm{loc}}(s, t)$ as the average over an ensemble of M independent simulation runs and over a population N for a given global opinion state s

$$
s^{\mathrm{loc}}(s, t) = \frac{1}{MN} \sum_M \sum_N s_i^{\mathrm{loc}}(s, t). \tag{4}
$$

$s^{\mathrm{loc}}(s, t)$ is the state of the neighborhood as it is perceived locally by an agent of any opinion including the particles' noisy perceptions and averaged over an ensemble of simulation runs. The relation between $s^{\mathrm{loc}}(s, t)$ and $s(t)$ was determined empirically and as expected it was found to be almost linear, see **Figure 1B**. Hence, we follow a two-step process of first accounting for the known influence by the agent's imperfect perception (γ) and then studying the remaining deviation. We approximate it linearly

$$
s^{\mathrm{loc}}(s, t) = c_1 s(t) + c_2. \tag{5}
$$

For perfect perception $\gamma = 0$ we would expect $s^{\mathrm{loc}}(s, t) \approx c_1 s(t)$. Here we have $\gamma = 0.2$. s^{loc} is time-variant but converges approximately to

$$
s^{\mathrm{loc}}(s, t') \approx (1 - 2\gamma)s(t') + \gamma, \tag{6}
$$

for $t' \gg 0$. This is a modeling approach based on a well-mixed assumption, not an attempt to fit the measured data. In

_____

[3]$s_1^* = 0$ and $s_2^* = 1$ are stable fixed points because in these states all particles have the same opinion and without noise these system states are absorbing (i.e., once reached the system cannot leave them anymore). Also note the instable fixed point $s_3 = 0.5$.

a well-mixed system (without any spatial correlations) an agent perceives a local state of $s^{loc} = \gamma$ for $s = 0$ and $s^{loc} = 1 - \gamma$ for $s = 1$. For $\gamma = 0.2$ we get $s^{loc}(s'_t) \approx 0.6s_{t'} + 0.2$ (plotted in **Figure 1B**). A plot of $g(s^{loc}(s_t), 0)$ in **Figure 1A** shows the effect of s^{loc}. The two stable fixed points move inwards toward the middle and the absolute values decrease.

$g(s^{loc}(s_t), a = 0)$ gives the exact values of Δs_t for an assumed well-mixed setting, that is, if an agent's current opinion is truly uncorrelated with its position. However, the spatial distribution of agents is not independent from the agents' opinion. A spatial correlation is likely to emerge based on the definition of the agent's behavior. The opinion of an agent is directly influenced by its neighbors, consequently they are correlated. This can also be determined in simulation by measuring the average, locally perceived opinion state for agents of a given opinion (measured during the last 100 time steps of the simulation, that is, $7,900 < t \leqslant 8,000$); $s_A^{loc}(s)$ gives the state perceived by agents of opinion A and $s_B^{loc}(s)$ gives the state perceived by agents of opinion B. The

differences between $s_{A,B}^{loc}(s)$ and $0.6s + 0.2$ show clearly a bias depending on the agents' current opinion as seen in **Figure 1C** which is evidence of a correlated spatial distribution of agents. The neighborhood of an agent with opinion A is populated by more agents of opinion A than on average for any state (for $s < 0.8$) and respectively for agents with opinion B.

We need an adjustment of the Galam model to account for spatial correlations. Although there is no explicit concept of contrarians in the density classification scenario, the contrarian approach can be used to compensate the effect of spatial correlations. That way the contrarians reflect the observed bias away from the current global majority due to local effects. For $s<0.5$ contrarians model the excess of perceived particles with opinion $o=A$ by particles with opinion A and for $s > 0.5$ contrarians model the excess of perceived particles with opinion $o=B$ by particles with opinion B. With increasing contrarian density a the two stable fixed points move further inwards until they would unite for $a \approx 0.0555$ leaving one stable

A three settings of the Galam model

B measured s^{loc} (circles), $0.6s + 0.2$ (line)

C spatial correlations in s_A^{loc} and s_B^{loc} based on an agent's opinion

D system dynamics Δs and fitted Galam model, $a \approx 0.025$, $d \approx 0.045$.
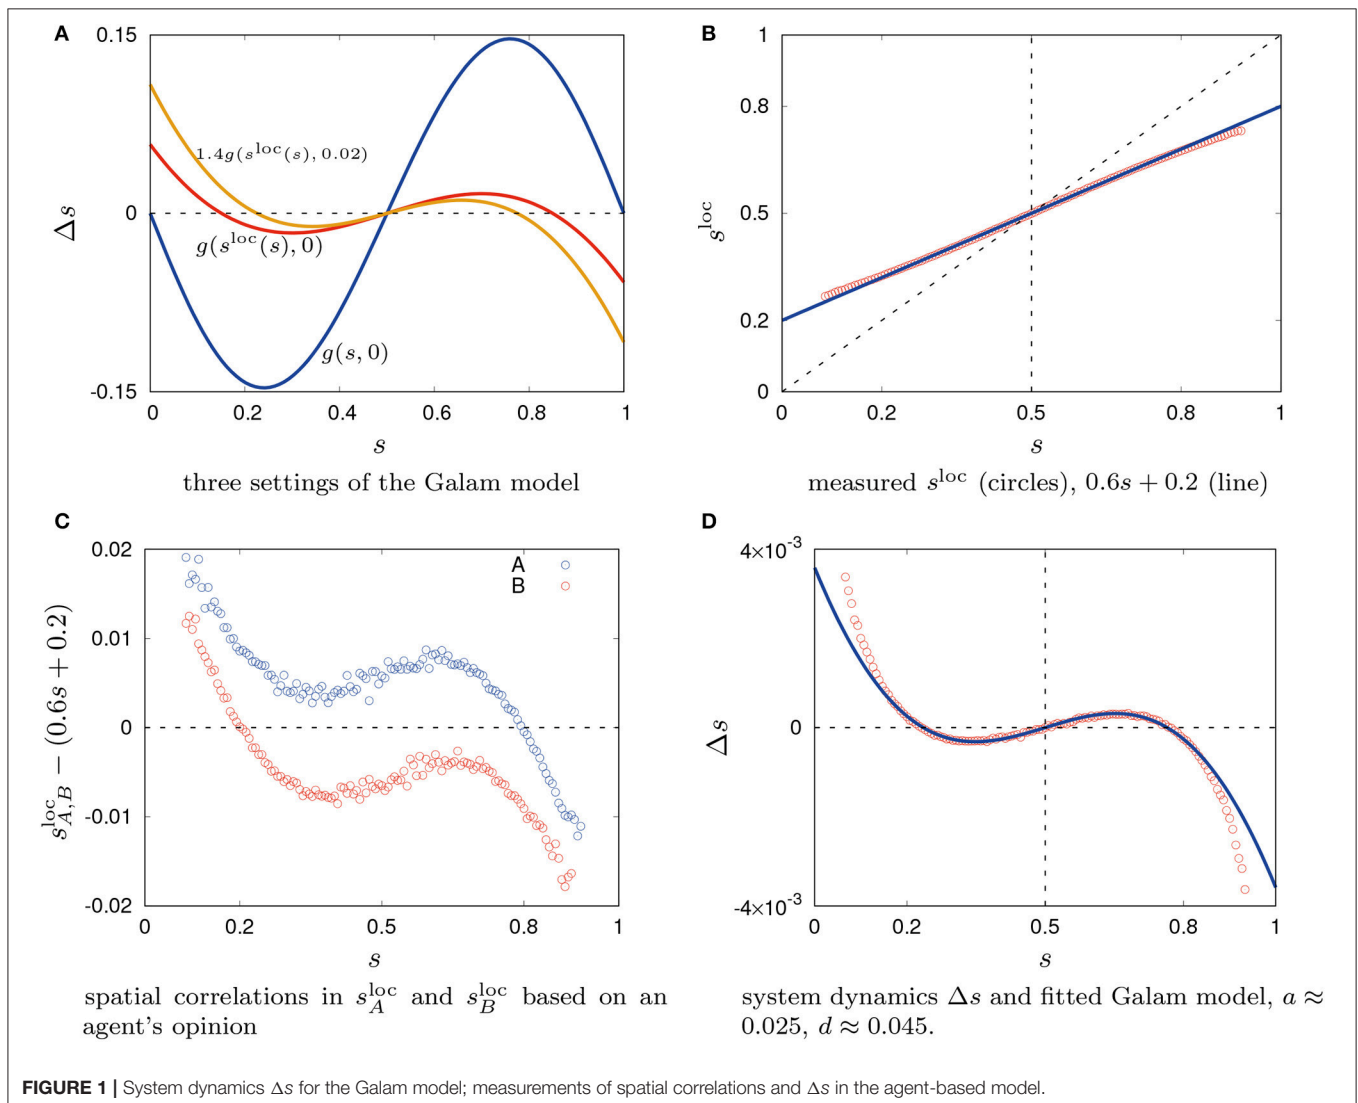
FIGURE 1 | System dynamics Δs for the Galam model; measurements of spatial correlations and Δs in the agent-based model.

fixed point at $s = 0.5$. In addition, we also scale the absolute values of $g(s^{\text{loc}}(s_t), a)$ by multiplying a constant d as done in **Figure 1A**. Fitting $dg(s^{\text{loc}}(s_t), a)$ via a and d to the empirically obtained data gives a good result for values $0.17 < s < 0.83$ but has systematic errors outside of that interval, see **Figure 1D**[4]. Data was obtained by measuring values of Δs_t as a function of the current system state s_t (i.e., $\Delta s_t(s_t) = s_{t+1} - s_t$) during the last 100 time steps of the simulation, that is, $7,900 < t \leqslant 8,000$. Plotted values are averages over all samples collected of the respective $\Delta s_t(s_t)$.

4. MODELING APPROACH II: URN MODEL

As an alternative to Galam's model we apply an urn model that was originally introduced as a model of collective decision-making in swarms (Hamann, 2012). The main idea is that we have a state-dependent probability of positive feedback $P_{\text{fb}}(s)$. The current majority opinion spreads for $P_{\text{fb}}(s) > 0.5$ and is diminished otherwise.

The idea of this urn model is as follows. An urn is filled with N agents which are either associated with opinion A or B. The game's dynamics is turn-based. First an agent is drawn with replacement and its opinion is noted. Then the opinion of a second agent is changed determined by that noted opinion. Say, first, an agent with opinion A is drawn. The probability of drawing an agent with opinion A is implicitly determined by the current number of agents with opinion A in the urn. The subsequent change of opinion of a second agent is determined by the probability of positive feedback $P_{\text{fb}}(s)$ and effects either a positive (an agent in the urn changes from opinion B to A, the fraction of the first drawn agent increases) or a negative feedback (an agent changes from opinion A to B, the fraction of the first drawn agent decreases). The feedback is determined explicitly by probability $P_{\text{fb}}(s)$ that we define below and that also depends on the current global opinion state s. Following Hamann (2012), the state variable's dynamics is defined by[5]
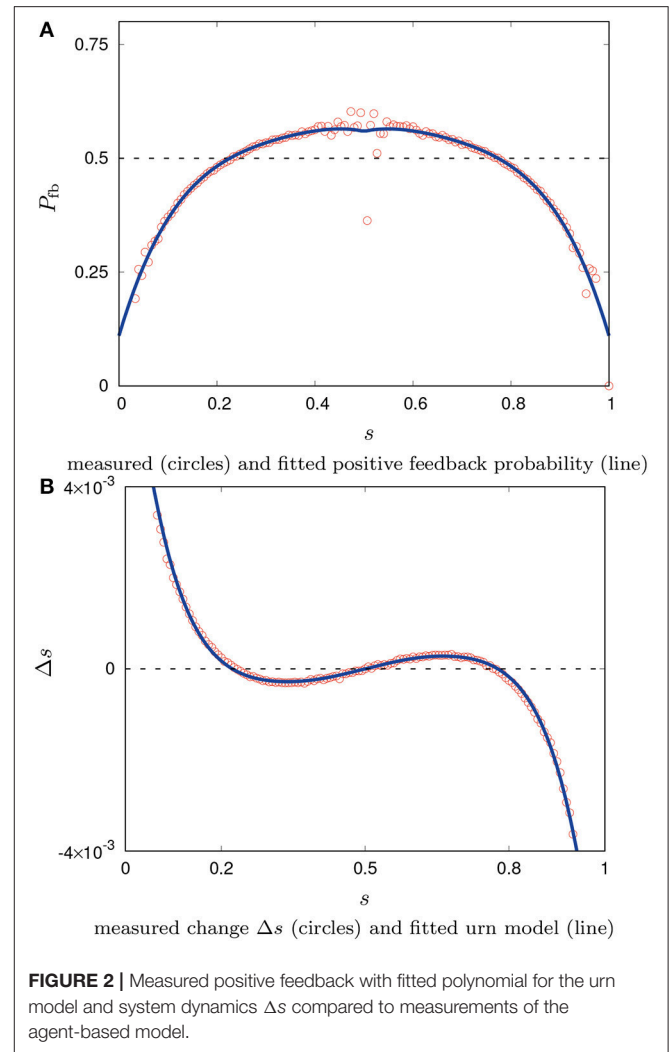
$$\Delta s_t = s_{t+1} - s_t = 4e(P_{\text{fb}}(s_t) - 0.5)(s_t - 0.5), \quad (7)$$

for a scaling constant e. The rationale of the urn model is to emulate, by the first draw, the frequency that an agent of a certain opinion happens to persuade another agent. The second draw models the average success rate of the persuasion based on the current global state. Thus, the urn model has no explicit concept of group sizes as Galam's model and only implicitly assumes a minimal setting of a bilateral meeting. Also spatial correlations of agents are not incorporated explicitly but can be represented by the probability of positive feedback $P_{\text{fb}}(s)$.

Following Hamann (2013), the probability of positive feedback can be measured in the simulation based on

FIGURE 2 | Measured positive feedback with fitted polynomial for the urn model and system dynamics Δs compared to measurements of the agent-based model.

observations of opinion revisions

$$P_{\text{fb}}(s) = \frac{\frac{r_b(s)}{r_b(s)+r_a(s)} - 1 + s}{2s - 1}, \text{ for } s \neq 0.5,$$

$$\min(s, 1-s) \leqslant \frac{r_b(s)}{r_b(s) + r_a(s)} \leqslant \max(s, 1-s), \quad (8)$$

for $r_b(s)$ is the absolute number of observed individual decision revisions from opinion A to B over any given period and $r_a(s)$ denotes revisions from B to A. The measured function $P_{\text{fb}}(s)$ is fitted by a polynomial of 4th degree[6] which is set mirror-symmetrical in $s = 0.5$. The result is shown in **Figure 2A**. Based on this empirically obtained function $P_{\text{fb}}(s)$ the dynamics of the system is then defined by Equation (7). A comparison to data from simulations is shown in **Figure 2B** which shows a very good fit[7].

---

[4]Summed squared error: 4.741×10^{-3}.

[5]This equation is easily obtained by basic probability theory considering all four cases of drawing an agent of either color followed by either positive or negative feedback.

[6]polynomial $f(x) = c_3 x^4 + c_4 x^3 + c_5 x^2 + c_6 x + c_7$ with $0 \leqslant f(x) \leqslant 1$ to model a probability, $c_3 = -21.3144$, $c_4 = 29.7651$, $c_5 = -16.1583$, $c_6 = 4.2788$, $c_7 = 0.0720295$.

[7]Summed squared error: 3.715×10^{-7}.

5. DISCUSSION AND CONCLUSION

We have reported two approaches to overcome the limitations of the well-mixed assumption in models of opinion dynamics. Both approaches are a combination of mathematical modeling and empirically obtained parameters.

The 2-state Galam opinion dynamics model with contrarians has systematic errors for extreme values of s, see **Figure 1D**. We chose to interpret the effect of spatial correlations as a contrarian effect. Hence, we used the main empirical element of Galam's model, parameter a specifying the fraction of contrarians, but the model still suffers from the simplifying assumption of well-mixed particles. The errors could also not be overcome by simple extensions of fitted s^{loc}-functions (data not shown). Despite this shortcoming, we gain a valid insight. A spatial correlation can be a local group of particles that share a similar opinion. However, they can be contrarian to the global majority. While this local group of particles acts according to properly defined decision rules, its global effect is that they oppose the majority as if they would defect the system in the way Galam's contrarians do.

In the case of the urn model, a very good fit to the simulation data was obtained using the probability of positive feedback that is measured following Equation (8). The urn model has a comprehensive empirical element $[P_{\text{fb}}(s)]$ and is still simple, concise, and achieves high accuracy. Measuring the positive feedback probability seems to comprise the averaged influence of correlations in the agents' spatial distribution. The urn model, hence, could be used to predict the long-term behavior of the collective system. The gained insight from this second modeling approach is that spatial correlations may be difficult to measure but they can be captured with a concise global modeling approach.

In addition, the knowledge about Δs can be used as a novel tool in multiple ways to model opinion dynamics in mobile agents. For example, one can macroscopically model the system dynamics as a Markov chain (Valentini et al., 2014, 2017) or by Langevin and Fokker–Planck equations (Carlen et al., 2013; Hamann, 2013) which allows for good predictions without modeling spatial distributions explicitly. Features such as the steady state of the probability density function of the global opinion state s or the mean first passage time (i.e., the switching time between two states of consensus) can be predicted with such models (Yates et al., 2009). Another interesting aspect is to apply the concept of the local opinion state $s^{\text{loc}}_{A,B}(s)$ within an agent to find an accurate estimate of the global state based on local sampling. One faces a kind of a bootstrapping problem then, because the agent only has a local sample instead of the actual global state. However, it seems feasible that systematic spatial correlations could be reduced by such a modeling approach. This would be useful especially in swarm robotics.

AUTHOR CONTRIBUTIONS

The author confirms being the sole contributor of this work and approved it for publication.

ACKNOWLEDGMENTS

REFERENCES

Carlen, E., Chatelin, R., Degond, P., and Wennberg, B. (2013). Kinetic hierarchy and propagation of chaos in biological swarm models. *Phys. D Nonlinear Phenomena* 260, 90–111. doi: 10.1016/j.physd.2012.05.013

Czirók, A. and Vicsek, T. (2000). Collective behavior of interacting self-propelled particles. *Phys. A* 281, 17–29. doi: 10.1016/S0378-4371(00)00013-3

Degond, P., and Yang, T. (2010). Diffusion in a continuum model of self-propelled particles with alignment interaction. *Math. Models Methods Appl. Sci.* 20, 1459–1490. doi: 10.1142/S0218202510004659

Franks, N. R., Dornhaus, A., Fitzsimmons, J. P., and Stevens, M. (2003). Speed versus accuracy in collective decision making. *Proc. R. Soc. Lond. B* 270, 2457–2463. doi: 10.1098/rspb.2003.2527

Galam, S. (1997). Rational group decision making: a random field Ising model at T=0. *Phys. A* 238, 66–80. doi: 10.1016/S0378-4371(96)00456-6

Galam, S. (2004). Contrarian deterministic effect on opinion dynamics: the "hung elections scenario". *Phys. A* 333, 453–460. doi: 10.1016/j.physa.2003.10.041

Galam, S. (2008). Sociophysics: a review of Galam models. *Int. J. Modern Phys. C* 19, 409–440. doi: 10.1142/S0129183108012297

Galam, S., and Jacobs, F. (2007). The role of inflexible minorities in the breaking of democratic opinion dynamics. *Phys. A* 381, 366–376. doi: 10.1016/j.physa.2007.03.034

Galam, S. and Moscovici, S. (1991). Towards a theory of collective phenomena: consensus and attitude changes in groups. *Eur. J. Soc. Psychol.* 21, 49–74. doi: 10.1002/ejsp.2420210105

Galstyan, A., Hogg, T., and Lerman, K. (2005). "Modeling and mathematical analysis of swarms of microscopic robots," in *Proceedings of IEEE Swarm Intelligence Symposium (SIS-2005), Pasadena, CA* (Los Alamitos, CA: IEEE Press), 201–208.

Garnier, S., Gautrais, J., Asadpour, M., Jost, C., and Theraulaz, G. (2009). Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adapt. Behav.* 17, 109–133. doi: 10.1177/1059712309103430

Gross, T., and Blasius, B. (2008). Adaptive coevolutionary networks: a review. *J. R. Soc. Interface* 5, 259–271. doi: 10.1098/rsif.2007.1229

Haken, H. (1977). *Synergetics - An Introduction*. Berlin: Springer.

Hamann, H. (2010). *Space-Time Continuous Models of Swarm Robotics Systems: Supporting Global-to-Local Programming*. Berlin: Springer.

Hamann, H. (2012). "Towards swarm calculus: universal properties of swarm performance and collective decisions," in *Swarm Intelligence: 8th International Conference, ANTS 2012*, Vol. 7461, LNCS, eds M. Dorigo, M. Birattari, C. Blum, A. L. Christensen, A. P. Engelbrecht, R. Groß, and T. Stützle (Berlin: Springer), 168–179.

Hamann, H. (2013). Towards swarm calculus: urn models of collective decisions and universal properties of swarm performance. *Swarm Intell.* 7, 145–172. doi: 10.1007/s11721-013-0080-0

Hamann, H. (2018). *Swarm Robotics: A Formal Approach*. Basel: Springer.

Hamann, H., Meyer, B., Schmickl, T., and Crailsheim, K. (2010). "A model of symmetry breaking in collective decision-making," in *From Animals to Animats 11*, Vol. 6226, LNAI, eds. S. Doncieux, B. Girard, A. Guillot, J. Hallam, J.-A. Meyer, and J.-B. Mouret (Berlin: Springer), 639–648.

Hamann, H., and Wörn, H. (2007). Embodied computation. *Parallel Process. Lett.* 17, 287–298. doi: 10.1142/S0129626407003022

Hegselmann, R., and Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Soc. Soc. Simul.* 5, 1–24. Available online at: http://jasss.soc.surrey.ac.uk/5/3/2.html

Helbing, D., and Molnar, P. (1997). "Self-organization phenomena in pedestrian crowds," in *Self-Organization of Complex Structures. From Individual to Collective Dynamics*, ed F. Schweitzer (London: Gordon and Breach). 569–577.

Helbing, D., Schweitzer, F., Keltsch, J., and Molnár, P. (1997). Active walker model for the formation of human and animal trail systems. *Phys. Rev. E* 56, 2527–2539. doi: 10.1103/PhysRevE.56.2527

Hillen, T., and Painter, K. J. (2009). A user's guide to PDE models for chemotaxis. *Math. Biol.* 58, 183–217. doi: 10.1007/s00285-008-0201-3

Holme, P., and Newman, M. E. J. (2006). Nonequilibrium phase transition in the coevolution of networks and opinions. *Phys. Rev. E* 74:056108. doi: 10.1103/PhysRevE.74.056108

Huepe, C., Zschaler, G., Do, A.-L., and Gross, T. (2011). Adaptive-network models of swarm dynamics. *New J. Phys.* 13:073022. doi: 10.1088/1367-2630/13/7/073022

Levine, H., Rappel, W.-J., and Cohen, I. (2000). Self-organization in systems of self-propelled particles. *Phys. Rev. E* 63:17101. doi: 10.1103/PhysRevE.63.017101

Mateo, D., Kuan, Y. K., and Bouffanais, R. (2017). Effect of correlations in swarms on collective response. *Sc. Rep.* 7:10388. doi: 10.1038/s41598-017-09830-w

Motsch, S., and Tadmor, E. (2014). Heterophilious dynamics enhances consensus. *SIAM Rev.* 56, 577–621. doi: 10.1137/120901866

Nicolis, S. C. and Dussutour, A. (2008). Self-organization, collective decision making and resource exploitation strategies in social insects. *Eur. Phys. J. B* 65, 379–385. doi: 10.1140/epjb/e2008-00334-3

Okubo, A. (1986). Dynamical aspects of animal grouping: swarms, schools, flocks, and herds. *Adv. Biophys.* 22, 1–94.

Packard, N. H. (1988). "Adaptation toward the edge of chaos," in *Dynamic Patterns in Complex Systems*, eds. J. A. S. Kelso, A. J. Mandell, and M. F. Shlesinger (Singapore: World Scientific), 293–301.

Prorok, A., Correll, N., and Martinoli, A. (2011). Multi-level spatial modeling for stochastic distributed robotic systems. *Int. J. Robot. Res.* 30, 574–589. doi: 10.1177/0278364911399521

Schimansky-Geier, L., Mieth, M., Rosé, H., and Malchow, H. (1995). Structure formation by active Brownian particles. *Phys. Lett. A* 207, 140–146. doi: 10.1016/0375-9601(95)00700-D

Schmickl, T. and Hamann, H. (2011). "BEECLUST: a swarm algorithm derived from honeybees," in *Bio-inspired Computing and Communication Networks*, ed Y. Xiao (Boca Raton, FL: CRC Press) 95–137.

Schmickl, T., Thenius, R., Möslinger, C., Radspieler, G., Kernbach, S., and Crailsheim, K. (2008). Get in touch: cooperative decision making based on robot-to-robot collisions. *Auton. Agents Multi Agent Syst.* 18, 133–155. doi: 10.1007/s10458-008-9058-5

Schweitzer, F. (2002). "Brownian agent models for swarm and chemotactic interaction," in *Fifth German Workshop on Artificial Life. Abstracting and Synthesizing the Principles of Living Systems*, eds D. Polani, J. Kim, and T. Martinetz (Akademische Verlagsgesellschaft Aka), 181–190.

Schweitzer, F. (2003). *Brownian Agents and Active Particles. On the Emergence of Complex Behavior in the Natural and Social Sciences.* Berlin: Springer.

Schweitzer, F., Zimmermann, J., and Mühlenbein, H. (2002). Coordination of decisions in a spatial agent model. *Phys. A* 303, 189–216. doi: 10.1016/S0378-4371(01)00486-1

Sood, V., and Redner, S. (2005). Voter model on heterogeneous graphs. *Phys. Rev. Lett.* 94:178701. doi: 10.1103/PhysRevLett.94.178701

Toner, J., and Tu, Y. (1998). Flocks, herds, and schools: a quantitative theory of flocking. *Phys. Rev. E* 58, 4828–4858. doi: 10.1103/PhysRevE.58.4828

Valentini, G., Ferrante, E., and Dorigo, M. (2017). The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Front. Robot. AI* 4:9. doi: 10.3389/frobt.2017.00009

Valentini, G., Hamann, H., and Dorigo, M. (2014). "Self-organized collective decision making: the weighted voter model," in *Proceedings of the 13th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, eds A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns (IFAAMAS), 45–52.

Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., and Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.* 6, 1226–1229. doi: 10.1103/PhysRevLett.75.1226

Vicsek, T., and Zafeiris, A. (2012). Collective motion. *Phys. Rep.* 517, 71–140. doi: 10.1016/j.physrep.2012.03.004

Yates, C. A., Erban, R., Escudero, C., Couzin, I. D., Buhl, J., Kevrekidis, I. G., et al. (2009). Inherent noise can facilitate coherence in collective swarm motion. *Proc. Natl. Acad. Sci. U.S.A.* 106, 5464–5469. doi: 10.1073/pnas.0811195106

frontiers
in Robotics and AI

Inform: Efficient Information-Theoretic Analysis of Collective Behaviors

Douglas G. Moore[1], Gabriele Valentini[1], Sara I. Walker[1]* and Michael Levin[2]

[1] BEYOND: Center for Fundamental Concepts in Science, Arizona Sate University, Tempe, AZ, United States, [2] Department of Biology, Allen Discovery Center, Tufts University, Medford, MA, United States

The study of collective behavior has traditionally relied on a variety of different methodological tools ranging from more theoretical methods such as population or game-theoretic models to empirical ones like Monte Carlo or multi-agent simulations. An approach that is increasingly being explored is the use of information theory as a methodological framework to study the flow of information and the statistical properties of collectives of interacting agents. While a few general purpose toolkits exist, most of the existing software for information theoretic analysis of collective systems is limited in scope. We introduce Inform, an open-source framework for efficient information theoretic analysis that exploits the computational power of a C library while simplifying its use through a variety of wrappers for common higher-level scripting languages. We focus on two such wrappers here: PyInform (Python) and rinform (R). Inform and its wrappers are cross-platform and general-purpose. They include classical information-theoretic measures, measures of information dynamics and information-based methods to study the statistical behavior of collective systems, and expose a lower-level API that allow users to construct measures of their own. We describe the architecture of the Inform framework, study its computational efficiency and use it to analyze three different case studies of collective behavior: biochemical information storage in regenerating planaria, nest-site selection in the ant *Temnothorax rugatulus*, and collective decision making in multi-agent simulations.

Keywords: information transfer, information storage, information dynamics, complex systems, collective behavior, information theory

1. INTRODUCTION

Collective behaviors, such as the coordinated motion of a flock of starlings (Ballerini et al., 2008), the collective decisions made by bees and ants (Franks et al., 2002), and the coordination of individual cells towards the creation or repair of a complex anatomical structure during embryogenesis or regeneration (Pezzulo and Levin, 2015), are complex collective phenomena that emerge from local interactions between many individuals. The study of these complex phenomena has been approached from many different angles, e.g., population models based on ordinary differential equations to predict the dynamics and study the stability of collective behaviors (Couzin et al., 2005; Marshall et al., 2009); game-theoretic approaches to study the emergence of cooperative strategies (Challet and Zhang, 1997); and multi-agents simulations to explore systems in the detail (Goldstone and Janssen, 2005). Another interesting approach is to focus on the distributed computation performed by the individuals in the collective (Langton, 1990; Mitchell, 1996; Lizier et al., 2014) and use information theory to analyze its architecture. Information theory has

been used, for example, to detect leadership relations between zebra fishes (Butail et al., 2016; Mwaffo et al., 2017) or to study foraging behavior of ant colonies (Reznikova and Ryabako, 1994; Zenil et al., 2015; Meyer, 2017). Additionally, it is extensively employed in the study of other complex systems with applications ranging from computational neuroscience (Honey et al., 2007; Vakorin et al., 2009; Lizier et al., 2011; Wibral et al., 2014), collectives of artificial agents (Williams and Beer, 2010; Boedecker et al., 2012; Walker et al., 2013; Biehl et al., 2016), neural and Boolean network models (Lizier et al., 2009; Kim et al., 2015; Walker et al., 2016), and multi-robot systems (Sperati et al., 2008; Sperati et al., 2011). Computing information theoretic measures, however, is computationally demanding and requires efficient software methodologies.

A common approach is to develop software solutions to compute specific information-theoretic measures. For example, TRENTOOL (Lindner et al., 2011) and MuTE (Montalto et al., 2014) are Matlab toolkits to compute transfer entropy. MVGC (Barnett and Seth, 2014) has been developed to compute Granger causality while ACSS (Gauvrit et al., 2016) and OACC (Soler-Toscano et al., 2014) to compute approximations to Kolmogorov complexity. However, while software options can always be developed to focus on particular techniques or methods, this approach is time-consuming for end-users. It can be tedious to explore and analyze the complex behavior of systems if every measure one chooses to use requires a separate library, not to mention the time spent in search of the functionality. What's more, it is not always easy to find a library to suit one's needs. One solution is to develop and make use of general-purpose software frameworks which can be applied across domains, and can provide researchers from different disciplines with a common software toolkit. At the risk of overselling our current endeavour, we can liken this approach to the development of solid, powerful linear algebra libraries such as BLAS (Lawson et al., 1979) and LAPACK (Anderson et al., 1999) which provide vast array of features and greatly simplify scientific computation. The most notable effort in this direction is the Java Information Dynamics Toolkit (JIDT) developed by (Lizier, 2014). JIDT is a Java library that provides access to classic information-theoretic measures (e.g., entropy and mutual information) as well as more recent measures of information dynamics (e.g., active information and transfer entropy) for both discrete and continuous data. JIDT is general-purpose and, thanks to the flexibility of the Java Virtual Machine, it can be called from several different high-level languages such as Matlab, Python or R.

In previous work (Moore et al., 2017), we introduced Inform: an open-source, general-purpose and cross-platform framework to perform information-theoretic analysis of collective of agents. Inform is a framework to analyze *discretely-valued*[1] time series data and is built to achieve two grounding objectives: computational efficiency and user flexibility. The first of these objectives is achieved by the core component of Inform, a high efficiency C library that takes care of the computation of information measures. The second objective is achieved through the design of a simple API and the development of a suite of wrappers for common higher-level programming languages, e.g., Python, R, Julia, and the Wolfram Language. The use of C as the implementation language and

Inform's carefully designed API make wrapping the core functionality straightforward. Since Inform has no external dependencies, distributing packages is greatly simplified. This is an advantage over libraries implemented in languages such as Java or R which require a virtual machine or an interpreter. Inform provides easy access to functions for empirically estimating probability distributions and uses them to compute common information-theoretic measures while also exposing a flexible API that a user can leverage to implement their own specialized measures. Additionally, Inform provides a collection of utilities that can be combined with other components of the framework to yield a wider range of analyses than those explicitly implemented. Inform provides a wide range of standard information-theoretic measures defined over time series and empirical probability distributions, as well as all of the common information dynamics measures. In addition, Inform provides a suite of functions for computing less common information-theoretic measures such as partial information decomposition (Williams and Beer, 2010), effective information (Hoel et al., 2013) and information flow (Ay and Polani, 2008). Inform v1.0.0 is released under the MIT license and is publicly available on GitHub[2].

In this work, we introduce two of Inform's language wrappers: PyInform[3] (Python) and rinform[4] (R). While the Inform library is, at least by C standards, straightforward to use, it is rather low-level. The decision to use C puts some of the memory-management burden on the user, and leads to rather rudimentary error handling. It is for these reasons that we invest the time in developing and maintain usable wrappers in a variety of higher-level languages. Without this initiative, users would have to call the C functions directly, decreasing the researcher's productivity and cluttering their code. This is not to mention the error-prone nature of interfacing languages. By targeting some of the more common languages used in the field, we aim to make the software and algorithms accessible to a wide user-base. The language wrappers are designed to provide users with an experience that is idiomatic to their chosen language under the assumption that users will be more productive in a language with which they are familiar. Inform's language wrappers are developed using the wrapping languages' native technology, e.g., object-orientation in Python. This allows users to work with a programming interface written in their chosen language without requiring knowledge of the core C library but still benefiting from its implementation of optimized algorithms.

We begin with a review of the design and implementation of the Inform framework in Section 2. In Section 2.1 we describe the architecture of Inform and its wrappers with a focus on each of the four major components of the framework—distributions, information measures, time series measures and utilities. In Section 2.2 we discuss the validation process and stability of Inform, PyInform and rinform. In Section 3 we showcase the capabilities of the framework by analyzing three different collective systems: cellular-level biochemical processes in regenerating planaria (see Section 3.1), house-hunting behavior in Temnothorax ants (see Section 3.2), and consensus achievement in multi-agent simulations (see Section 3.3). Section 4 is dedicated to the analysis of the computational performance

---

[1]*While the current release of Inform only supports analysis of discrete time series, full support for continuous data is planned, see Section 6.*

of Inform taking the JIDT library of (Lizier, 2014) as the reference framework and using active information and transfer entropy as benchmark metrics. Section 5 presents demonstrative examples of how to use PyInform and rinform with simple use cases for each of Inform v1.0.0's major components. Finally, Section 6 concludes this paper with a discussion of the advantages and the shortcomings of the Inform framework as well as a summary of future directions of development.

2. DESIGN AND IMPLEMENTATION

Inform (MIT license)[5] is a general-purpose library and framework for information-theoretic analysis of empirical time series data. Much of the design of Inform has focused on making the library (and its language wrappers) as intuitive and easy to use as possible, all the while attempting to provide powerful features that *some* other toolkits lack. Some of Inform's features include:

- Optimized implementations of many common information-theoretic time series measures, including block entropy, mutual information, complete and apparent transfer entropy, active information storage and predictive information.
- Optimized implementations of less common concepts such as effective information, information flow, evidence for integration and partial information decomposition.
- All time series measures include local and average variants where applicable.
- An empirical probability distribution structure over a discrete event space[6] and a suite of basic information-theoretic functions built around it.
- A collection of utility functions, such as black boxing and binning algorithms, which may be used in conjunction with time series measures to facilitate analysis of complex systems.
- No external library dependencies.

The Inform library is implemented in cross-platform C, and can be built on any system with a C11-compliant[7] compiler. The choice of C was not a simple one. The decision came down to two factors:

1. Essentially all modern programming languages provide a C foreign-function interface.
2. Most of Inform's functionality requires minimal memory management — typically only one allocation and deallocation per function.
3. C does not have exceptions. While useful in a given language, exceptions make interfacing languages more difficult.
4. C requires no external dependencies for distribution — as such, the wrapper libraries do not depend on an external virtual machine, interpreter or JIT compiler.

---

[5]https://github.com/elife-asu/inform
[6]Support for continuous event spaces is planned for v2.0.0, Section 6.
[7]ISO/IEC 9899:2011: https://www.iso.org/standard/57853.html

All subsequent references to Inform will refer to the entire framework including its wrappers; any reference to the C library will be disambiguated as such.

2.1. Architecture

Information theory largely focuses on quantifying information within probability distributions. To model this, Inform is designed around the concept of an empirical probability distribution. These distributions are used to define functions which compute information theoretic quantities. From these basic building blocks, we implemented an entire host of time series measures. Intuitively, the time series measures construct empirical distributions and call the appropriate information-theoretic functions. These three components—distributions, information measures and time series measures—form Inform's core functionality. Additionally, Inform provides a suite of utilities that can be used to augment and extend it's core features. We now detail how these components are implemented and interact with each other to provide a cohesive toolkit.

Inform's empirical probability distributions are implemented by a distribution class, `Dist`. This class, which is a wrapper for the C structure `inform_dist`, stores the relative frequencies of observed events that can then be used to estimate each event's probability. The framework provides a suite of functions built around `Dist` which makes it easy for users to create distributions, accumulate observations and output probability estimates. It is important to note that Inform's empirical distributions are only defined for discrete events. Subsequent releases will natively support continuous data (see Section 6).

Inform uses the `Dist` class to provide well-defined implementations of many Shannon information measures. In Python, the canonical example of such a function is

```
pyinform.shannon.entropy(dist, b = 2)
```

which computes the (Shannon) entropy of the distribution `dist` using a base-`b` logarithm . Equivalently, the R function to compute Shannon entropy is given by

```
shannon_entropy(dist, b = 2)
```

Each measure in the framework takes some number of distributions and the logarithmic base as arguments, ensures that they are all valid[8], and returns the desired quantity. Inform v1.0.0 only provides information measures based on Shannon's notion of entropy, but other types are planned for future releases (see Section 6).

Inform's final core component is a suite of measures defined over time series. The version 1.0.0 release includes 15 time series measures with average and local (sometimes referred to as pointwise) variants provided where applicable. Each measure essentially performs some variation on the same basic procedure: first, accumulate observations from the time series into empirical distributions, and then, use them to compute some distribution-based information measure. **Table 1** provides a complete list of the time series measures provided in Inform v1.0.0.

---

[8]An empirical distribution is considered invalid if it has no recorded events.

TABLE 1 | The time series measures available in inform v1.0.0.

Time Series Measure	Local/Pointwise Variant
Block Entropy (Shannon, 1948)	✓
Cross Entropy (Cover and Thomas, 2005)	✗*
(Multivariate) Mutual Information (Tononi et al., 1994; Cover and Thomas, 2005)	✓
Conditional Entropy (Cover and Thomas, 2005)	✓
Relative Entropy (Kullback and Leibler, 1951; Cover and Thomas, 2005)	✓
Entropy Rate (Cover and Thomas, 2005)	✓
Active Information (Lizier et al., 2012)	✓
Transfer Entropy (Schreiber, 2000; Kaiser and Schreiber, 2002; Lizier et al., 2008)	✓
Separable Information (Lizier et al., 2010)	✓
Predictive Information (Bialek et al., 2001a; Bialek et al., 2001b)	✓
Excess Information (Crutchfield and Feldman, 2003; Feldman and Crutchfield, 2003)	✓
Effective Information (Hoel et al., 2013; Hoel, 2017)	✗
Information Flow (Ay and Polani, 2008)	✗
Partial Information Decomposition (Williams and Beer, 2010)	✗
Evidence of Integration (Biehl et al., 2016)	✗

Local/Pointwise variants are implemented for all measures that reasonably admit them, signified by a ✓ . A ✗ denotes measures for which a local variant is not implemented.
*(✗) Cross entropy's local variant is equivalent to local block entropy, and is thus not implemented.

The final component of Inform is the utility suite. One of the greatest challenges of building a general-purpose framework is ensuring that it can be applied to problems that are outside of the authors' initial use cases. Inform attempts to do this by first exposing the basic components of the library, distributions and information measures, and then providing utility functions that can be used to augment the core functionality. One particular example of this is the `black_box`[9] function which losslessly produces a single time series from a collection of time series (see Section 5.4 for a detailed description and an example of use of this particularly versatile function). The `black_box` function allows Inform to avoid implementing multivariate variants of time series measures while still making it straightforward for users to compute such quantities. Of course, there are a multitude of uses for such a function. Our aim is that the utility suite can extend Inform's functionality well beyond what the authors had in mind when implementing the core library.

2.2. Validation

The Inform framework was developed using a test-driven approach: unit tests were written for each component before implementing the component itself. Consequently, all features in Inform have been thoroughly unit tested to ensure that they perform as expected. In fact, the bulk of the development effort went into testing, and test code accounts for roughly 60% of the entire C source code distribution.

To ensure cross-platform support, continuous integration services are employed to build and run all unit tests on multiple platforms. Travis CI[10] builds currently ensure support for Linux with the gcc 4.6.3 and clang 3.4 compilers, and Mac OS X with AppleClang 7.3.0.7030031. AppVeyor[11] builds ensure support for Windows with Microsoft Visual Studio 14 2015. Code coverage reports for PyInform and rinform are hosted by CodeCov[12] and currently show a coverage of 97% and 91%, respectively, while coverage for the C implementation is in the works for future releases.

3. ANALYSIS OF COLLECTIVE BEHAVIORS

In this section, we illustrate the use of Inform by performing information-theoretic analyses of three collective behaviors: the dynamics membrane potentials and ion concentrations in regenerating planaria, nest-site selection by colonies of the ant *Temnothorax rugatulus*, and collective decision-making in a multi-agent system. While the following results are interesting in their own right, and will likely be considered more deeply in subsequent work, our primary focus is on showcasing the utility and range of the Inform framework.

3.1. Biochemical Collectivity in Regenerating Planaria

In this first case study, we use partial information decomposition (Williams and Beer, 2010) to analyze how various ions contribute to the cell membrane potentials in a regenerating planarian. Planaria are an order of flatworms which have prodigious regenerative abilities (Sheĭman and Kreshchenko, 2015). When a planarian is cut in half, each piece will regenerate the missing tissue and develop into a fully functional individual. Recent work is stored in a complex biophysical circuit which is not hardwired by the genome (Oviedo et al., 2010; Beane et al., 2011; Emmons-Bell et al., 2015; Durant et al., 2017). Many pharmacological reagents that target the endogenous bioelectrical machinery (ion channels and electrical synapses known as gap junctions) can alter the behavior of this circuit and thus alter the large-scale bodyplan to which fragments regenerate. An example of this is ivermectin, a chloride channel opener, which results in the development of a two-headed phenotype upon regeneration (Beane et al., 2011). The resulting two-headed morphology is persistent under subsequent regeneration events outside of the presence of ivermectin. The hypothesis is that these gap-junction inhibitors disrupt proper bio-electric communication between cells and lead the organism to non-wildtype morphological attractors. As an initial step at understanding how the morphological information is stored and modified, we can look at how information about the bio-electric

---

[9]The naming of this function is intended to bring to mind the process of "black boxing" nodes in a network. That is, this function models drawing an opaque box around a collection of nodes, treating them as one unit with no known internal structure.

FIGURE 1 | The V_{mem} distribution over the body of a BETSE simulated planarian (Pietak and Levin, 2016) over the course of regeneration: **(A)** pre-surgery, **(B)** 0 s post-surgery, **(C)** 1000s post surgery. **(D)** The non-zero redundancy sub-lattice computed via partial information decomposition. Each node presents the redundant information provided by the given collection of random variables. Of the 166 nodes in the full redundancy lattice, these 13 are the only nodes which yield non-zero unique information. All other nodes were pruned, and the edges were constructed using the Williams-Beer dependency relations. Nodes are colored roughly by the order of magnitude of their unique information content.

patterning is stored in specific intracellular ion concentrations of Na^+, K^+, Ca^{2+} and Cl^-.

We use the BioElectric Tissue Simulation Engine (BETSE) (Pietak and Levin, 2016) to simulate the planarian regeneration process under a simple two-cut intervention (Pietak and Levin, 2017). For this demonstrative case study, we simulate the planarian for 1000 s after two surgical cuts are made, dividing the worm into three pieces **Figure 1A-C**. From the simulation we extract the time series, sampled at a frequency of $10Hz$ ($10,000$ time steps), of the average cell membrane potentials V_{mem} and the Na^+, K^+, Ca^{2+} and Cl^- ion concentrations for each cell. We use a "threshold" binning to bin the average cell membrane potentials using a biologically realistic activation threshold of $-40mV$, the cell is considered depolarized (state 1) when V_{mem} is above $-40mV$, and hyperpolarized (state 0) otherwise. Each of the ion concentrations are separately binned into two uniform bins whose sizes depend on the range of the ion's concentration.

From these binned data, we compute the partial information decomposition (PID) of the information about V_{mem} provided by the ion concentrations. From the 4 ion variables, Inform constructs the full 166-node redundancy lattice; however, only 13 of those nodes represent variable combinations that contribute unique information, in the sense of (Williams and Beer, 2010). We pruned all but those 13 variable combinations. The resulting sub-lattice is depicted in **Figure 1D**. Altogether, the intracellular ion concentrations yield approximately 0.425 bits of information about the average cell membrane potential – computed as the sum of the unique information provided by each node. This is less than the theoretical maximum of 1 bits, but that's hardly surprising given that the cell membrane potential is determined by the difference between the intra- and extracellular ion concentrations. We also see that the only individual ion that provides any unique information about V_{mem} is Na^+ – Na^+ is the only ion that appears alone in **Figure 1D**. We know that both Na^+ and K^+ play a crucial role

in determining V_{mem}, so it is surprising to see that Na^+ is the dominate information provider. Subsequent work will delve deeper into the what this decomposition tells us about the biochemical mechanisms of regeneration.

As we conclude this example, it is worthwhile to acknowledge that Inform's current implementation of PID is limited to Williams's and Beer's I_{min} measure of redundant information (Williams and Beer, 2010). A number of alternative measures of redundancy and uniqueness could be applied to the redundancy lattice, e.g. (Bertschinger et al., 2014), and there is continuing discussion as to which is the "correct" measure. A subsequent version of PID will allow the user to specify which measure they would prefer, and even allow them to implement their own.

3.2. Nest-Site Selection by the Ant *Temnothorax Rugatulus*

In this case study, we use local active information to analyze collective decisions made by the ant *Temnothorax rugatulus* (Pratt et al., 2002; Sasaki et al., 2013). Specifically, we consider nest-site selection, a popular and well-studied collective behavior observed both in honeybee swarms and ant colonies (Franks et al., 2002). When Temnothorax ants need to choose a new nest, individuals in the colony explore the surrounding environment looking for possible candidate sites (e.g., a rock crevice). Upon the identification of a good candidate, an ant may perform a tandem run—a type of recruitment process whereby the ant returns to the old nest to lead another member of the colony in a tandem to the newly found site for a possible assessment. Tandem runs, together with independent discoveries of the same site, allow for a build up of a population of ants at that site which in turn triggers the achievement of a quorum, i.e., the identification by individual ants of the popularity of a candidate site. After quorum is reached, ants switch from performing tandem runs to performing transport—a type of recruitment process distinct from tandem runs whereby an ant returns to the old nest, loads another ant on her back and carries that ant to a site. The combination of parallel exploration, tandem runs, quorum sensing and transports allows Temnothorax ants to concurrently evaluate different candidate sites and converge on a collective decision for the best one.

For this study, we look at a live colony of 78 T. *rugatulus* ants repeatedly choosing between a good and a mediocre site in a laboratory environment for a total of 5 experiments. We consider ants to be in one of three state: uncommitted (state 0), committed to the good site (state 1) or committed to the mediocre site (state 2). All ants in the colony are individually paint-marked using a four-color code which allows us to identify individual ants and track their commitment state. From video-recordings of the experiments, we extract the commitment state of each ant over time as follows: initially, all ants are considered uncommitted, and ants commit to a certain site after performing a tandem run or a transport towards that site or when they are transported to that site. We record the commitment state of each ant every second and obtain 78 time series for each of the 5 experiments which we use to compute the local active information (history length $k = 2$). As different experiments differ in duration due to the stochasticity inherent to colony emigrations, time series extracted
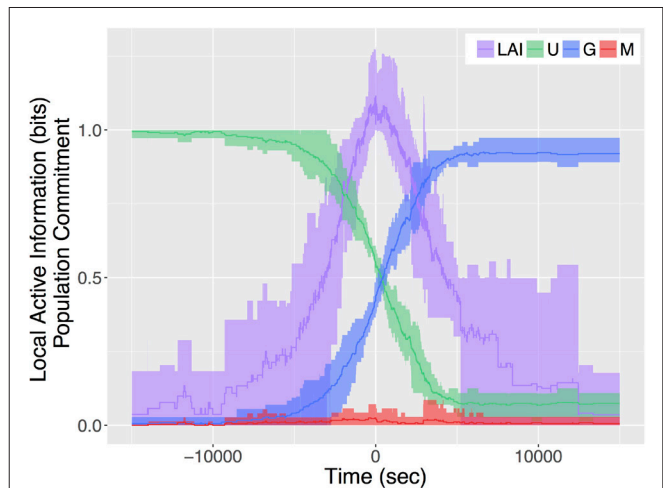

FIGURE 2 | Distribution of local active information and colony-level commitment state for a live colony of 78 T. *rugatulus* ants computed over 5 colony emigrations. Lines represent mean values of local active information (**LAI**), and proportions of ants in the colony that are uncommitted (**U**), committed to the good site (**G**) and committed to the mediocre site (**M**). Shaded areas correspond to minimum and maximum values of the same quantities.

from different experiments also differ in length (but all 78 time series within the same experiment have the same length). In our analysis, we considered shortened time series of 3×10^4 time steps (approximately the same duration of the fastest emigration experiment) following a procedure described below.

Figure 2 shows the results of our analysis of the local active information together with the change of commitment over time for the entire colony. Data are aggregated as follows: we first compute the mean local active information of individual ants in a colony emigration; then, we find the point in time where local active information peaks; finally, we center the local active information and the colony-level commitment state for each emigration around this point in time (i.e., time 0 in **Figure 2**) and compute mean, maximum and minimum values over experiments. The peak in the local active information is approximately in the middle of the decision-making process (i.e., when half of the colony is committed for the good site and half is still uncommitted). This maximum of the local active information, approximately 1 bit, identifies a critical point in the collective decision.

3.3. Multi-Agent Simulations

In this final case study, we use transfer entropy to analyze the flow of information in a multi-agent system developed to study the best-of-n problem (Valentini et al., 2017). Specifically, we consider a system where a collective of agents needs to chose between two options: 0 or 1. The behavior of each agent is defined as a probabilistic finite-state machine with 2 states for each option: exploration and dissemination. In the exploration state, an agent explores the environment and evaluates the quality of its currently favored option. In the dissemination state, an agent promotes its opinion (i.e., broadcast its preference for a particular option to its neighbors) for a time proportional to the quality of its favored option. At the

end of the dissemination state, soon before transitioning to the exploration state, the agent collects the preferences of its neighbors and applies a decision rule to reconsider its current preference. In this case study we consider two decision rules: the majority rule, whereby an agent adopts the option favored by the majority of its neighbors, and the voter model, whereby an agent adopts the option favored by a randomly chosen neighbor (Valentini et al., 2016).

We consider a collective of 100 agents tasked with a binary decision-making problem where the best option has quality 1.0 and the other option has quality 0.9. All the agents in the collective apply the same decision rule (i.e., either the majority rule or the voter model) over a neighborhood represented by the agent's 5 nearest-neighbors. For each decision rule, we performed 1000 multi-agent simulations where the initial preferences of the agents are equally distributed among the two options. We let simulations run for a total duration each of 10^4 seconds. Our aim is to use transfer entropy to analyze the flow of information to an agent from its neighborhood as it applies its decision rule. We extract a binary-state series of preferences for each agent, where each element of the series is the agent's preference immediately prior to apply it's decision rule. We then construct a 6-state series of neighborhood states, each element of which is the number of neighbors with a preference for the best option (i.e., $\{0, \ldots, 5\}$) at the time of the agent's decision. As opposed to the previous case study, each simulation lasts for the same amount of time. However, the number of applications of a decision rule by an agent within the same simulation and across different simulations is stochastic. Consequently, time series derived from different agents differ in length (on average, 13.93 ± 3.27 for the majority rule and 13.82 ± 2.89 for the voter model). To mitigate the effect of short time series, we used time series from all agents within a simulation to compute the probability distributions required for transfer entropy (i.e., an average of 1393 samples for the majority rule and 1382 for the voter model) and consider this quantity an average over all agents of the collective. In this system, agents are memoryless and parameters have been tuned to approximate a well-mixed interaction pattern. However, time correlation may still be present as a result of the interaction of agents with their neighborhood. For simplicity, we use a history length of $k = 1$ and let the investigation of longer history lengths for future work.

Figure 3 shows the results of our analyses of the multi-agent simulations. Specifically, it depicts the probability density functions (PDF) of the average transfer entropy toward an agent applying a decision rule over 1000 simulations. To compute the average transfer entropy towards an agent, we estimate the required probability distributions from the time series of all agents in the collective and use these distributions to obtain one sample of transfer entropy for each simulation. The PDFs of transfer entropy obtained for the majority rule and for the voter model are remarkably different (two sample t-test, p-value $< 2.2 \cdot 10^{-16}$). On average, the majority rule has a higher value of transfer entropy (0.3106 bits) with respect to the voter model (0.2019 bits). However, it is also characterized by a larger spread with a SD of 0.1302 bits compared to that of the voter model, 0.0301 bits. Previous analysis of these decision mechanisms under similar conditions showed that the majority rule is much faster than the voter model and its consensus time has an higher variance as well (Valentini et al., 2016). These results

FIGURE 3 | Probability density functions of the average transfer entropy for agents in systems applying the majority rule (purple) and for agents in systems using the voter model (green).
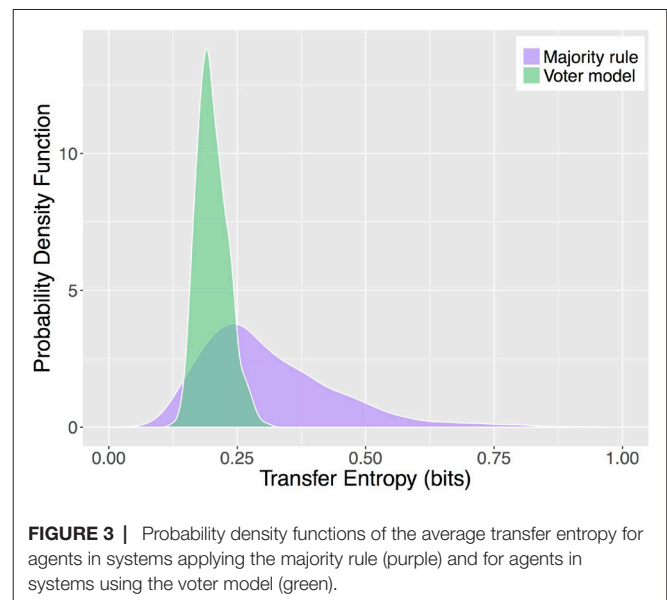
are likely correlated and a deeper analysis of this case study is currently undergoing.

4. PERFORMANCE ANALYSIS

In this section, we investigate the performance of PyInform by calculating two computationally demanding measures of information dynamics: active information (AI) and transfer entropy (TE). While we focus on PyInform here, rinform shows comparable performance characteristics. We compare the performance of PyInform with that of JIDT (Lizier, 2014) which we take as the gold-standard for the field. We chose AI and TE as they are the primary overlap in the functionality of PyInform and JIDT. The time series for the following tests were generated using the same multi-agent simulation described in Section 3.3. The state of each agent includes its opinion (i.e., 0 or 1) and its control state (i.e., dissemination or exploration). As such, the time series for each agent is base-4 and runs for the entire duration of the simulation, not just the decision points as in Section 3.3. We considered four different data sets wherein we varied both the decision rule (i.e., majority rule or voter model) and the difficulty of the decision-making problem (i.e., $\rho_0 = 1.0$ and $\rho_1 \in \{0.5, 0.9\}$). For each data set, we executed 1000 simulations with a duration of 1001 time steps using a collective of 50 agents initialized with an equal distribution of preferences for both options.

Using the four data sets described above, we computed the AI for each agent in the collective and the TE using PyInform and JIDT's built-in time series-based functionality. We computed AI and TE for history lengths $1 \leq k \leq 11$ or until computational resources were exhausted. For each data set and history length k, we repeated 5 times the calculations and timed the computational process. In computing the run times, we considered only the time necessary to loop over the agent combinations and to compute the relevant values while we disregarded the time spent reading data files and comparing results. All performance tests were single-threaded and
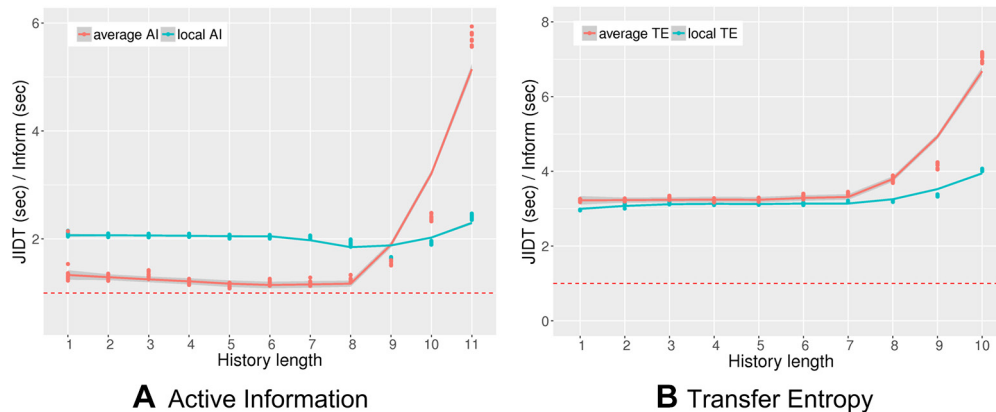
FIGURE 4 | Performance ratio versus history length for average and local active information **(A)** and average and local transfer entropy **(B)** The dashed lines mark a performance ratio of 1.0. Memory constraints limited computation of transfer entropy with JIDT up to $k = 10$.

run with Amazon Web Services, using a c4.large EC2 instance relying on a 2 vCPUs and 3.75 GB of RAM[13]

Figure 4 shows the results of the performance comparison as the ratio of execution times between JIDT and PyInform for active information (left panel) and transfer entropy (right panel). In both experiments, the PyInform package outperforms JIDT with a speedup ranging from a minimum of $1.2\times$ up to a maximum speedup of $7\times$. The computational gain of PyInform over JIDT is more pronounced when computing average measures with a history length $k > 8$ both in the case of AI and in that of TE. It is obligatory to note that history lengths $k > 8$ are rarely useful in practice as the amount of data necessary for the measures to show statistical significance grows exponentially in k. We include the longer history lengths, simply to acknowledge that both frameworks experience exponential growths in runtime as k grows. As one would expect, the computational requirements of transfer entropy are greater than those of active information for both frameworks.

In addition to comparing the runtime performance, we also compared the absolute results of the calculations for all values of k. The values computed with the PyInform package never differed from those of the JIDT library by more than 10^{-6} bits. PyInform is *marginally* more computationally efficient than JIDT while providing equally accurate calculations of information-theoretic measures. However, it is important to remember that computational performance is not the only aspect that one should consider when choosing a software solution. Developer time is often more valuable than computation time. For example, JIDT offers many benefits over Inform including its support for continuously-valued data and a wider range of parameters (e.g., source embedding, embedding delays, source-target delay). Subsequent versions of Inform will reduce the discrepancy in features (see Section 6), and the library wrappers are designed

to increase programmer productivity. Whether or not speed is a deciding factor in a user's decision to use Inform will depend on the requirements of the task at hand.

5. USE CASE EXAMPLES

In this section we provide a few examples of how to directly use the Python and R wrappers, respectively, PyInform and rinform. Live documentation of these wrappers can be found at https://elife-asu.github.io/PyInform and https://elife-asu.github.io/rinform.

5.1. Empirical Distributions

We start with a simple example of how to use the `Dist` class to estimate a probability distribution from a binary sequence of events (see **Listing 1** for PyInform and **Listing 2** for rinform). In Python, the `from_data` static method creates a distribution and records observations from an array of discrete events. The same objective can be achieved in R using the `infer` function. In this case, two observations are made of the event "0" and three of event "1". The `probability` method can be used to query the estimated probability of a given event. Alternatively, the `dump` method can then be used to return an array of all estimated probabilities.

Listing 1 | Estimate a probability distribution from a binary sequence of events. (Python)

```
In [1]: from pyinform import Dist
In [2]: dist = Dist.from_data([0,1,1,0,1]) # observe 2 0's and 3 1's
In [3]: dist
Out[3]: Dist.from_hist([2, 3])
In [4]: dist.probability(0) # What is the probability of seeing a 0?
Out[4]: 0.4
In [5]: dist.probability(1) # What is the probability of seeing a 1?
Out[5]: 0.6
In [6]: dist.dump() # output the probabilities to an array Out[6]: array([0.4, 0.6])
```

---

[13]See https://aws.amazon.com/ec2/instance-types/ for the specifications of the c4.large EC2 instance..

Listing 2 | **Estimate a probability distribution from a binary sequence of events. (R)**

```
In [1]: library(rinform)
In [2]: dist <- infer(c(0,1,1,0,1)) # observe 2 0's and 3 1's
In [3]: dist
Out[3]: $histogram: [1] 2 3
Out[3]: $size: [1] 2
Out[3]: $counts: [1] 5
Out[3]: attr(,"class"): [1] "Dist"
In [4]: probability(dist, 1) # What is the probability of seeing a 0?
Out[4]: 0.4
In [5]: probability(dist, 2) # What is the probability of seeing a 1?
Out[5]: 0.6
In [6]: dump(dist) # output the probabilities to an array
Out[6]: [1] 0.4 0.6
```

Listing 3 | **Estimate the entropy of an empirical distribution of binary events. (Python)**

```
In [1]: from pyinform import shannon
In [2]: from pyinform import Dist
In [3]: dist = Dist(2) # create a Dist over two events
In [4]: dist.accumulate([0,1,1,0,1]) # accumulate some observations
Out[4]: 5 # 5 observations were made
In [5]: shannon.entropy(dist, b = 2) # compute the base-2 Shannon entropy
Out[5]: 0.9709505944546686
```

This is only a sample of the functionality provided around the `Dist` class. Further examples can be found in the live documentation of PyInform[14] and rinform [15]

5.2. Shannon Information Measures

As described in Section 2.1, the Shannon information measures are defined around the `Dist` class. In this subsection, we give an example of how to compute the Shannon entropy of a distribution. In **Listing 3**, we demonstrate how to construct a `Dist` instance and compute its entropy using PyInform while **Listing 4** shows the equivalent implementation using rinform. The resulting distribution can record observations of two events, "0" or "1". With the distribution in hand, the `accumulate` function accumulates the observations from an array. This is functionally equivalent to `Dist.from_data` which was used in **Listing 1** (Python) and `infer` which was used in **Listing 2** (R). Once the distribution has been created, computing its entropy is as simple as performing a single function call to `shannon.entropy` (in Python) or `shannon_entropy` (in R).

A host of information measures are provided in the Inform framework. These can be found in the `pyinform.shannon` module[16] for PyInform. While rinform is not organized into modules, the user has access to all the same information measures as described in the rinform's documentation[17]

5.3. Time Series Measures

The time series measures are a primary focus for the Inform framework. **Listing 5** (Python) and **Listing 6** (R) provide a

[14]http://elife-asu.github.io/PyInform/dist.html
[15]https://elife-asu.github.io/rinform/#2_empirical_distributions.
[16]http://elife-asu.github.io/PyInform/shannon.html
[17]https://elife-asu.github.io/rinform/#3_shannon_information_measures

Listing 4 | **Estimate the entropy of an empirical distribution of binary events. (R)**

```
In [1]: library(rinform)
In [2]: dist <- Dist(2) # create a Dist over two events
In [3]: dist <- accumulate(dist, c(0,1,1,0,1)) # accumulate some observations
In [4]: shannon_entropy(dist, b = 2) # compute the base-2 Shannon entropy
Out[5]: [1] 0.9709506
```

Listing 5 | **Estimate the average and local transfer entropy from discrete data. (Python)**

```
In [1]: import numpy as np
In [2]: from pyinform import transfer_entropy
In [3]: np.random.seed(2018)
In [4]: src = np.random.randint(0, 4, 100)
In [5]: target = np.zeros(len(source), dtype = int)
In [6]: target[1:] =src[:–1]
In [7]: transfer_entropy(src, target, k = 1) # TE with history length 1
Out[7]: 1.8705725949309469
In [8]: lte = transfer_entropy(src, target, k = 1, local = True) # Local TE
In [9]: lte.shape
Out[9]: (1, 99)
In [10]: np.mean(lte) # the mean local TE is approximately the
Out[10]: 1.870572594930947 # same as Out[7]
In [11]: lte = transfer_entropy(src, target, k = 0)
        … # stack trace removed for brevity
InformError: an inform error occurred - "history length is zero"
```

Listing 6 | **Estimate the average and local transfer entropy from discrete data. (R)**

```
In [1]: library(rinform)
In [2]: set.seed(2018)
In [3]: src <- sample(0:3, 100, TRUE)
In [4]: target <- c(src[100], src[1:99])
In [5]: transfer_entropy(src, target, k = 1) # TE with history length 1
Out[5]: [1] 1.912181
In [6]: lte <- transfer_entropy(src, target, k = 1, local = TRUE) # Local TE
In [7]: dim(lte)
Out[7]: (99, 1)
In [8]: mean(lte) # the mean local TE is approximately the
Out[8]: [1] 1.912181 # same as Out[5]
In [9]: lte <- transfer_entropy(src, target, k = 0)
Out[9]: Error: <k > is less then 1!
```

complete example of how to estimate the average and local (pointwise) transfer entropy between two base-4 time series — this functionality was used in the performance analysis described in Section 4. To demonstrate this, we construct[18] a source time series, `src`, and then shift and copy it to a target time series, `target`. The expected result is that the average transfer entropy from `src` to `target` will be near 2.0 bits. The `transfer_entropy` function is employed to compute this value. The examples go on to compute the local transfer entropy, which returns an array of local (pointwise) values.

Time series measures can fail for a variety of reasons ranging from invalid arguments to exhausted system memory. In these situations, an error is raised which describes the reason for the

[18]In Python, we use —numpy—, a package that provides a wealth of useful array-based functionality: http://www.numpy.org/.

Listing 7 | Estimate the average multivariate active information of two continuous time series. (Python)

```
In [1]: from pyinform import active_info
In [2]: from pyinform.utils import bin_series, black_box
In [4]: threshold = 0.5
In [5]: node1, _, _ =bin_series([0.5, 0.2, 0.6, 0.8, 0.7], bounds = [threshold])
In [6]: node1
Out[6]: array([1, 0, 1, 1, 1], dtype = int32)
In [7]: node2, _, _ =bin_series([0.1, 0.9, 0.4, 0.7, 0.4], bounds = [threshold])
In [8]: node2
Out[8]: array([0, 1, 0, 1, 0], dtype = int32)
In [9]: series = black_box((node1, node2))
In [10]: series
Out[10]: array([2, 1, 2, 3, 2], dtype = int32)
In [11]: active_info(series, k = 1)
Out[11]: 1.
```

Listing 8 | Estimate the average multivariate active information of two continuous time series. (R)

```
In [1]: library(rinform)
In [3]: threshold <- 0.5
In [5]: node1 <- bin_series(c(0.5, 0.2, 0.6, 0.8, 0.7), bounds = threshold)$binned
In [6]: node1
Out[6]: [1] 1 0 1 1 1
In [7]: node2 <- bin_series(c(0.1, 0.9, 0.4, 0.7, 0.4), bounds = threshold)$binned
In [8]: node2
Out[8]: [1] 0 1 0 1 0
In [9]: series <- black_box(matrix(c(node1, node2), ncol = 2), l = 2)
In [10]: series
Out[10]: [1] 2 1 2 3 2
In [11]: active_info(series, k = 1)
Out[11]: [1] 1
```

function's failure. At the end of both **Listing 5** and **Listing 6**, we provide an example of an erroneous function invocation. Pyinform raises an `InformError` while rinform prints an error message.

All of the time series measures follow the same basic calling conventions as `transfer_entropy`. Further examples of the various time series measures can be found in the live documentation of PyInform[19] and rinform [20]

5.4. Utility Functions

Our next example, **Listing 7** and **Listing 8**, demonstrates how to use Inform's utility functions to estimate the multivariate active information of two continuous time series, `node1` and `node2`. It begins by binning points in each time series into one of two bins, $x < 0.5$ or $x \geq 0.5$, using the `bin_series` function. Once binned, the series are black-boxed, that is, their states are aggregated together over a larger state-space, using the `black_box` function to produce a base-4 time series (i.e., the product of the bases of `node1` and `node2`). Each time step of this black-boxed time series, `series`, represents the joint state of the two binned time series. From `series`, the multivariate active information with $k = 1$ is estimated using the `active_info` function.

The flexibility of the the `black_box` function makes it worthwhile to elaborate further on precisely what it does. In making concurrent observations of a collection of random variables, say $X_1, X_2, \ldots$, which may or may not be correlated with one another, we are in fact making observations of an underlying variable W defined over a different state space Ω. These observed variables can be thought of as views, filters or projections of the the underlying system state drawn from Ω. Many information analyses require the reconstruction of Ω from the observations of $X_1, X_2, \ldots$. The `black_box` function covers this role in Inform. Given a number of time series, each representing the time series of a random variable, `black_box` losslessly encodes the joint state of those time series as a single value in the system's joint state space Ω. As a concrete example, consider the following time series of concurrent observations of two random variables

$$X : \ 0, 1, 1, 0, 1, 0, 0, 1,$$
$$Y : \ 1, 0, 0, 2, 1, 2, 1, 2.$$

Here, X is a binary variable while Y is a trinary one. Together, observations of X and Y may be thought to represent observations of an underlying state variable $W = (X, Y) \in \Omega$[21]:

$$W : \ (0, 1), (1, 0), (1, 0), (0, 2), (1, 1), (0, 2), (0, 1), (1, 2).$$

As such, these observations can be encoded as a base-6 time series which is precisely what `black_box` does, yielding

$$W : \ 1, 3, 3, 2, 4, 2, 1, 5.$$

The `black_box` function accepts a host of arguments which augment how it constructs the resulting time series, all of which are described and demonstrated in the documentation[22].

Inform's collection of utilities allows the user to easily construct new information-measures over time series data. Combining utility functions such as `black_box` with common time series measures such as `mutual_info` is a powerful way for the user to extend the functionality of the Inform framework to include measures of particular interest to their research.

We will now conclude this section with two demonstrative examples of how `black_box` can be combined with the time series functions `block_entropy`[23] and `mutual_info` to implement *conditional entropy* and *active information*, respectively. First recall that the conditional entropy of a random variable X conditioned on a random variable Y is defined as

$$H(X \mid Y) = - \sum_{x,y} p(x, y) \log p(x \mid y) = H(X, Y) - H(Y). \quad (1)$$

As such, one might compute the conditional entropy by first constructing the joint distribution (X, Y) (using `black_box`) and

---

[19]http://elife-asu.github.io/PyInform/timeseries.html
[20]https://elife-asu.github.io/rinform/#4_time_series_measures.

[21]Note that if we had considered W = (Y,X) ∈ Ω instead, the encoded time series would have been different , e.g., 2,1,1,4,3,4,2,5. However, the mutual information between them, I(W,W), tends to the theoretical maximum H(W) as the number of observations increases; this indicates that (X,Y) and (Y,X) are informationally equivalent representations of the underlying space.
[22]http://elife-asu.github.io/PyInform/utils.html , https://elife-asu.github.io/rinform/#5_utilities.
[23]The —block_entropy— function computes the Shannon block entropy of a time series. This reduces to the standard Shannon entropy when a block size of k = 1 is used, e.g., —block_entropy(series, k = 1)—.

Listing 9 | Estimate conditional entropy between two time series using `black_box` **and** `block_entropy`**. (Python)**

```
In [1]: from pyinform import block_entropy, conditional_entropy
In [2]: from pyinform.utils import black_box
In [3]: X = [0,1,2,2,2,2,0,1,0] # the target variable
In [4]: Y = [0,0,1,1,1,1,0,0,0] # the condition variable
In [5]: XY = black_box((X,Y)) # the joint variable (X,Y)
In [6]: conditional_entropy(X, Y) # H(X | Y) =H(X,Y) - H(Y)
Out[6]: 0.5394169969192604
In [7]: block_entropy(XY, k = 1) - block_entropy(Y, k = 1)
Out[7]: 0.5394169969192604
```

Listing 11 | Estimate active information of a time series using `black_box` **and** `mutual_info`**. (Python)**

```
In [1]: from pyinform import active_info, mutual_info
In [2]: from pyinform.utils import black_box
In [3]: X = [0,0,1,1,1,1,0,0,0]
In [4]: X2 = black_box(X, k = 2) # the 2-histories of X
In [5]: active_info(X, k = 2)
Out[5]: 0.3059584928680418
In [6]: mutual_info(X[2:], X2[:−1]) # align indices of X and X2
Out[6]: 0.3059584928680421
```

then computing the difference of entropies as in Equation (1) (using `block_entropy`). This is demonstrated using PyInform in **Listing 9** and rinform in **Listing 10**.

Finally, we will perform a similar process to estimate the active information of random variable X as defined by

$$A_k(X) = \sum_{x^+, x^{(k)}} p(x^+, x^{(k)}) \log \frac{p(x^+, x_i^{(k)})}{p(x^+)p(x^{(k)})} = I(X^+, X^{(k)}) \quad (2)$$

where X^+ is the random variable representing the state of X in the next time step and $X^{(k)}$ is the present k-history of X. We can use `black_box` to construct the time series of k-histories, and `mutual_info` to compute the mutual information between X^+ and $X^{(k)}$ as in Equation (2). We demonstrate this using PyInform and rinform in **Listing 11** and **Listing 12**, respectively.

6. CONCLUSION AND DISCUSSION

In this paper we introduced Inform v1.0.0, a flexible and computationally efficient framework to perform information-theoretic analysis of collective behaviors. Inform is a general-purpose, open-source, and cross-platform framework designed to be flexible and easy to use. It builds on a computationally efficient C library and an ecosystem of foreign language wrappers for Python, R, Julia, and the Wolfram Language. Inform gives the user access to a large set of functions to estimate information-theoretic measures from empirical discretely-valued time series. These include classic information-theoretic measures such as Shannon's entropy and mutual information, information dynamics measures such as active information storage and transfer entropy, and information-based concepts conceived to investigate the causal architecture of collective systems. Inform's low-level API

is organized around the concepts of probability distributions, information measures, time series measures and utilities and its flexibility allows users to construct new measures and algorithms of their own. We showcased the Inform framework by applying it to the study of three collective behaviors: cellular-level biochemical processes in regenerating planaria, colony emigration by the ant *Temnothorax rugatulus*, and collective decision-making in multi-agent simulations. We investigated the performance of the Inform framework by comparing them with those of the JIDT library showing that Inform have similar or superior performance with respect to JIDT. In effect, Inform is a potentially invaluable tool for any researcher performing information analysis of collective behaviors and other complex systems.

The Inform framework is still a relatively young project compared to more mature projects such as JIDT. While it has many features that make it unique such as, its computational efficiency, the large set of information-theoretic methods, and the availability of foreign language wrappers, it does lack some important functionality. We are planning three subsequent releases to incrementally extend the Inform framework. In the version 1.1.0 release, we will modify Inform's interface to provide the user with access to the probability distributions used in the computation of information dynamics measures and their accumulation functions. In Python, for example, the extended API for computing the active information may take the following form:

```
class ActiveInfoAccumulator(Accumulator):
  def __init__(self):
    pass
  def accumulate(self, data):
    pass
  def evaluate(self, local = False):
    pass
```

The advantage of exposing probability distributions and their accumulation functions is that the user can modify the way

Listing 10 | Estimate conditional entropy between two time series using `black_box` **and** `block_entropy`**. (R)**

```
In [1]: library(rinform)
In [2]: X <- c(0, 1, 2, 2, 2, 2, 0, 1, 0) # the target variable
In [3]: Y <- c(0, 0, 1, 1, 1, 1, 0, 0, 0) # the condition variable
In [4]: XY <- black_box(matrix(c(X, Y), ncol = 2), l = 2) # the joint variable (X,Y)
In [5]: conditional_entropy(Y, X) # H(X | Y) =H(X,Y) - H(Y)
Out[5]: 0.539417
In [6]: block_entropy(XY, k = 1) - block_entropy(Y, k = 1)
Out[6]: 0.539417
```

Listing 12 | Estimate active information of a time series using `black_box` **and** `mutual_info`**. (R)**

```
In [1]: library(rinform)
In [3]: X <- c(0, 0, 1, 1, 1, 1, 0, 0, 0)
In [4]: X2 <- black_box(X, l = 1, r = 2) # the 2-histories of X
In [5]: active_info(X, k = 2)
Out[5]: 0.3059585
In [6]: mutual_info(matrix(c(X[3:9], X2[1:7]), ncol = 2))
Out[6]: 0.3059585
```

that probabilities are estimated. As opposed to the version 1.0.0 where Inform's time series measures require that all time series be stored in memory prior to the estimation of distributions, this new release will allow the user to write their own accumulation functions which could incrementally update distributions from very large time series stored on the hard-drive or with data that is generated in real-time. In the version 1.2.0 release, we will provide support for non-Shannon entropy functions. Shannon's entropy of a discrete random variable is the unique functional form of entropy that satisfies all Shannon's four axioms (Shannon, 1948). However, many functional forms of entropy become possible as soon as these four axioms are relaxed or otherwise modified. Two examples of such non-Shannon entropy forms are Rényi entropy (Rényi, 1961) and Tsallis-Havrda-Charvát entropy (Havrda and Charvát, 1967; Tsallis, 1988). Shannon's entropy is currently used in the calculations of most information dynamics measures available in Inform. The version 1.2.0 release will allow the user to make use of Non-Shannon entropy functions which may give insight into the dynamics of information processing in non-ergodic systems. Finally, the version 2.0.0 release will represent a major improvement of the Inform framework by providing support for continuously-valued time series. Although Inform provides utilities to discretize continuous data through the process of binning, its repertoire of information-theoretic measures only supports discretely-valued time series. Discretely-valued time series allows for computational efficiency (complexity is $O(N)$ in the length of the time series N), however, the discretization of continuous data might introduce artifacts and reduce the accuracy of the overall analysis. In the version 2.0.0 release we will implement estimation techniques for continuous probability distributions, such as kernel density estimation (Rosenblatt, 1956; Parzen, 1962; Schreiber, 2000; Kaiser and Schreiber, 2002), with the aim of extending Inform's reach towards continuously-valued data. More advanced estimation techniques, such as Kraskov-Stögbauer-Grassberger estimation (Kraskov et al., 2004), are planned for subsequent releases once we have a standardized API support of continuous data. Some additional details concerning future releases of the Inform framework are described on the Issues page[24] of the GitHub repository where users are encouraged to suggest features or report bugs.

AUTHOR CONTRIBUTIONS

DGM designed and implemented the Inform library as well as the Python, Julia, and Mathematica wrappers. GV designed and implemented the R wrapper. All authors contributed to the conceptualization of the framework and to the writing of the manuscript.

FUNDING

ACKNOWLEDGMENTS

---

[24]https://github.com/elife-asu/inform/issues

REFERENCES

Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J, J., and Du Croz, J. (1999). *LAPACK users' guide*, Third Edn. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Ay, N., and Polani, D. (2008). Information flows in causal networks. *Advs. Complex Syst.* 11 (01), 17–41. doi: 10.1142/S0219525908001465

Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., et al. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proc. Natl. Acad. Sci. U.S.A.* 105 (4), 1232–7. doi: 10.1073/pnas.0711437105

Barnett, L., and Seth, A. K. (2014). The MVGC multivariate Granger causality toolbox: A new approach to Granger-causal inference. *J. Neurosci. Methods* 223, 50–68. doi: 10.1016/j.jneumeth.2013.10.018

Beane, W. S., Morokuma, J., Adams, D. S., and Levin, M. (2011). A chemical genetics approach reveals H,K-ATPase-Mediated Membrane membrane voltage is required for planarian head regeneration. *Chem. Biol.* 18 (1), 77–89. doi: 10.1016/j.chembiol.2010.11.012

Bertschinger, N., Rauh, J., Olbrich, E., Jost, J., and Ay, N. (2014). Quantifying unique information. *Entropy* 16 (4), 2161–83. doi: 10.3390/e16042161

Bialek, W., Nemenman, I., and Tishby, N. (2001a). Predictability, complexity, and learning. *Neural Comput.* 13 (11), 2409–63. doi: 10.1162/089976601753195969

Bialek, W., Nemenman, I., and Tishby, N. (2001b). Complexity through nonextensivity. *Physica A* 302 (1-4), 89–99. doi: 10.1016/S0378-4371(01)00444-7

Biehl, M., Ikegami, T., and Polani, D. (2016). "Towards information based spatiotemporal patterns as a foundation for agent representation in dynamical systems" *Proceedings of the Artificial Life Conference 2016* (Cambridge, MA), pp. 722–729.

Boedecker, J., Obst, O., Lizier, J. T., Mayer, N. M., and Asada, M. (2012). Information processing in echo state networks at the edge of chaos. *Theory Biosci.* 131 (3), 205–13. doi: 10.1007/s12064-011-0146-8

Butail, S., Mwaffo, V., and Porfiri, M. (2016). Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Phys. Rev. E* 93 (4):042411. doi: 10.1103/PhysRevE.93.042411

Challet, D., and Zhang, Y. -C. (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A* 246 (3-4), 407–18. doi: 10.1016/S0378-4371(97)00419-6

Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature* 433 (7025), 513–6. doi: 10.1038/nature03236

Cover, T. M., and Thomas, J. A. (2005). *Elements of Information Theory*. Hoboken, NJ, USA: John Wiley & Sons, Inc.

Crutchfield, J. P., and Feldman, D. P. (2003). Regularities unseen, randomness observed: Levels of entropy convergence. *Chaos* 13 (1), 25–54. doi: 10.1063/1.1530990

Durant, F., Morokuma, J., Fields, C., Williams, K., Adams, D. S., and Levin, M. (2017). Long-term, stochastic editing of regenerative anatomy via targeting endogenous bioelectric gradients. *Biophys. J.* 112 (10), 2231–43. doi: 10.1016/j.bpj.2017.04.011

Emmons-Bell, M., Durant, F., Hammelman, J., Bessonov, N., Volpert, V., Morokuma, J., et al. (2015). Gap junctional blockade stochastically induces different species-specific head anatomies in genetically wild-type girardia dorotocephala flatworms. *Int. J. Mol. Sci.* 16 (11), 27865–96. doi: 10.3390/ijms161126065

Feldman, D. P., and Crutchfield, J. P. (2003). Structural information in two-dimensional patterns: Entropy convergence and excess entropy. *Phys. Rev. E* 67 (5):051104. doi: 10.1103/PhysRevE.67.051104

Franks, N. R., Pratt, S. C., Mallon, E. B., Britton, N. F., and Sumpter, D. J. T. (2002). Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society B: Biological Sciences* 357 (1427), 1567–83. doi: 10.1098/rstb.2002.1066

Gauvrit, N., Singmann, H., Soler-Toscano, F., and Zenil, H. (2016). Algorithmic complexity for psychology: a user-friendly implementation of the coding theorem method. *Behav. Res. Methods* 48 (1), 314–29. doi: 10.3758/s13428-015-0574-3

Goldstone, R. L., and Janssen, M. A. (2005). Computational models of collective behavior. *Trends Cogn. Sci. (Regul. Ed.).* 9 (9), 424–30. doi: 10.1016/j.tics.2005.07.009

Havrda, J., and Charvát, F. (1967). Quantification method of classification processes. concept of structural -entropy. *Kybernetika* 3 (1), 30–5.

Hoel, E. (2017). When the map is better than the territory. *Entropy* 19 (5):188. doi: 10.3390/e19050188

Hoel, E. P., Albantakis, L., and Tononi, G. (2013). Quantifying causal emergence shows that macro can beat micro. *Proc. Natl. Acad. Sci. U.S.A.* 110 (49), 19790–5. doi: 10.1073/pnas.1314922110

Honey, C. J., Kötter, R., Breakspear, M., and Sporns, O. (2007). Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proc. Natl. Acad. Sci. U.S.A.* 104 (24), 10240–5. doi: 10.1073/pnas.0701519104

Kaiser, A., and Schreiber, T. (2002). Information transfer in continuous processes. *Physica D: Nonlinear Phenomena* 166 (1-2), 43–62. doi: 10.1016/S0167-2789(02)00432-3

Kim, H., Davies, P., and Walker, S. I. (2015). New scaling relation for information transfer in biological networks. *J. R. Soc. Interface* 12 (113):20150944. doi: 10.1098/rsif.2015.0944

Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. *Phys. Rev. E* 69 (6):066138. doi: 10.1103/PhysRevE.69.066138

Kullback, S., and Leibler, R. A. (1951). On information and sufficiency. *Ann. Math. Statist.* 22 (1), 79–86. doi: 10.1214/aoms/1177729694

Langton, C. G. (1990). Computation at the edge of chaos: phase transitions and emergent computation. *Physica D: Nonlinear Phenomena* 42 (1-3), 12–37. doi: 10.1016/0167-2789(90)90064-V

Lawson, C. L., Hanson, R. J., Kincaid, D. R., and Krogh, F. T. (1979). Basic linear algebra subprograms for fortran usage. *ACM Trans. Math. Softw.* 5 (3), 308–23. doi: 10.1145/355841.355847

Lindner, M., Vicente, R., Priesemann, V., and Wibral, M. (2011). TRENTOOL: a matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC Neurosci.* 12 (1):119. doi: 10.1186/1471-2202-12-119

Lizier, J. T. (2014). JIDT: an information-theoretic toolkit for studying the dynamics of complex systems. *Front. Robot. AI* 1 (December), 1–20. doi: 10.3389/frobt.2014.00011

Lizier, J. T., Heinzle, J., Horstmann, A., Haynes, J. -D., and Prokopenko, M. (2011). Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *J. Comput. Neurosci.* 30 (1), 85–107. doi: 10.1007/s10827-010-0271-2

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2008). Local information transfer as a spatiotemporal filter for complex systems. *Phys. Rev. E* 77 (2):026110. doi: 10.1103/PhysRevE.77.026110

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2010). Information modification and particle collisions in distributed computation. *Chaos* 20 (3), 037109. doi: 10.1063/1.3486801

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2012). Local measures of information storage in complex distributed computation. *Inf. Sci. (Ny).* 208, 39–54. doi: 10.1016/j.ins.2012.04.016

Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2014). A framework for the local information dynamics of distributed computation in complex systems. In Guided self-organization: inception. 115–158.

Lizier, JT., Piraveenan, M., Pradhana, D., Prokopenko., and Yaeger, LS. (2009). "Functional and structural topologies in evolved neural networks" *European Conference on Artificial Life* pp. 140–7.

Marshall, J. A., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., and Franks, N. R. (2009). On optimal decision-making in brains and social insect colonies. *J. R. Soc. Interface* 6 (40), 1065–74. doi: 10.1098/rsif.2008.0511

Meyer, B. (2017). Optimal information transfer and stochastic resonance in collective decision making. *Swarm Intell.* 11 (2), 131–154. doi: 10.1007/s11721-017-0136-7

Mitchell, M. (1996). Computation in cellular automata: a selected review. *Nonstandard Computation* 95–140.

Montalto, A., Faes, L., and Marinazzo, D. (2014). MuTE: a matlab toolbox to compare established and novel estimators of the multivariate transfer entropy. *PLoS ONE* 9 (10):e109462. doi: 10.1371/journal.pone.0109462

Moore, DG., Valentini, G., Walker, SI., and Levin, M. (2017). "Inform: a toolkit for information-theoretic analysis of complex systems" *Proceedings of the 2017 IEEE Symposium Series on Computational Intelligence* (Honolulu, HI, USA), 1–8.

Mwaffo, V., Butail, S., and Porfiri, M. (2017). Analysis of pairwise interactions in a maximum likelihood sense to identify leaders in a group. *Front. Robot. AI* 4:35. doi: 10.3389/frobt.2017.00035

Oviedo, N. J., Morokuma, J., Walentek, P., Kema, I. P., Gu, M. B., Ahn, J. -M., et al. (2010). Long-range neural and gap junction protein-mediated cues control polarity during planarian regeneration. *Dev. Biol.* 339 (1), 188–199. doi: 10.1016/j.ydbio.2009.12.012

Parzen, E. (1962). On estimation of a probability density function and mode. *Ann. Math. Statist.* 33 (3), 1065–1076. doi: 10.1214/aoms/1177704472

Pezzulo, G., and Levin, M. (2015). Re-membering the body: applications of computational neuroscience to the top-down control of regeneration of limbs and other complex organs. *Integr. Biol.* 7 (12), 1487–1517. doi: 10.1039/C5IB00221D

Pietak, A., and Levin, M. (2016). Exploring instructive physiological signaling with the bioelectric tissue simulation engine. *Front. Bioeng. Biotechnol.* 4 (Pt 4), 55. doi: 10.3389/fbioe.2016.00055

Pietak, A., and Levin, M. (2017). Bioelectric gene and reaction networks: computational modelling of genetic, biochemical and bioelectrical dynamics in pattern regulation. *J. R. Soc. Interface* 14 (134):20170425. doi: 10.1098/rsif.2017.0425

Pratt, S., Mallon, E., Sumpter, D., and Franks, N. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albipennis*. *Behav. Ecol. Sociobiol. (Print).* 52 (2), 117–127. doi: 10.1007/s00265-002-0487-x

Rényi, A. (1961). "On measures of entropy and information" *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics* pp. 547–561.

Reznikova, Z., and Ryabako, B. (1994). Experimental study of the ants' communication system with the application of the information theory approach. *Memorabilia Zoologica* 48, 219–236.

Rosenblatt, M. (1956). Remarks on some nonparametric estimates of a density function. *Ann. Math. Statist.* 27 (3), 832–837. doi: 10.1214/aoms/1177728190

Sasaki, T., Granovskiy, B., Mann, R. P., Sumpter, D. J. T., and Pratt, S. C. (2013). Ant colonies outperform individuals when a sensory discrimination task is difficult but not when it is easy. *Proc. Natl. Acad. Sci. U.S.A.* 110 (34), 13769–13773. doi: 10.1073/pnas.1304917110

Schreiber, T. (2000). Measuring Information Transfer. *Phys. Rev. Lett.* 85 (2), 461–464. doi: 10.1103/PhysRevLett.85.461

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal* 1928, 379–423.

Sheĭman, I. M., and Kreshchenko, N. D. (2015). Regeneration of planarians: experimental object. *Russ. J. Dev. Biol.* 46 (1), 1–9. doi: 10.1134/S1062360415010075

Soler-Toscano, F., Zenil, H., Delahaye, J. -P., Gauvrit, N, Nicolas, G. (2014). Calculating kolmogorov complexity from the output frequency distributions of small turing machines. *PLoS ONE* 9 (5):e96223. doi: 10.1371/journal.pone. 0096223

Sperati, V., Trianni, V., and Nolfi, S. (2008). Evolving coordinated group behaviours through maximisation of mean mutual information. *Swarm Intell.* 2 (2-4), 73–95. doi: 10.1007/s11721-008-0017-1

Sperati, V., Trianni, V., and Nolfi, S. (2011). Self-organised path formation in a swarm of robots. *Swarm Intell.* 5 (2), 97–119. doi: 10.1007/s11721-011-0055-y

Tononi, G., Sporns, O., and Edelman, G. M. (1994). A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proc. Natl. Acad. Sci. U.S.A.* 91 (11), 5033–5037. doi: 10.1073/pnas.91.11.5033

Tsallis, C. (1988). Possible generalization of Boltzmann-Gibbs statistics. *J. Stat. Phys.* 52 (1-2), 479–487. doi: 10.1007/BF01016429

Vakorin, V. A., Krakovska, O. A., and McIntosh, A. R. (2009). Confounding effects of indirect connections on causality estimation. *J. Neurosci. Methods* 184 (1), 152–160. doi: 10.1016/j.jneumeth.2009.07.014

Valentini, G., Ferrante, E., and Dorigo, M. (2017). The best-of-*n* problem in robot swarms: Formalization, state of the art, and novel perspectives. *Front. Robot. AI* 4:9. doi: 10.3389/frobt.2017.00009

Valentini, G., Ferrante, E., Hamann, H., and Dorigo, M. (2016). Collective decision with 100 Kilobots: speed versus accuracy in binary discrimination problems. *Auton. Agent. Multi. Agent. Syst.* 30 (3), 553–580. doi: 10.1007/s10458-015-9323-3

Walker, SI., Cisneros, L., and Davies, P. (2013). "Evolutionary transitions and top-down causation" *Proceedings of Artificial Life XIII* 283–290.

Walker, SI., Kim, H., and Davies, P. C. (2016). The informational architecture of the cell. *Philos. Trans. A Math. Phys. Eng. Sci.* 374 (2063):20150057. doi: 10.1098/rsta.2015.0057

Wibral, M., Vicente, R., and Lizier, J. T. (2014). *Directed information measures in neuroscience*. Germany: Springer.

Williams, P., and Beer, R. (2010). Information dynamics of evolved agents. *From Animals to Animats* 11, 38–49.

Williams, P. L., and Beer, R. D. (2010). *Nonnegative Decomposition of Multivariate Information*. New York: Cornell University, 1–14.

Zenil, H., Marshall, JAR., and Tegnér, J. (2015). "*Approximations of algorithmic and structural complexity validate cognitive-behavioural experimental results*" *Quantitative Methods (q-bio.QM)* 1–33.

frontiers
in Robotics and AI

Virtual Sensing and Virtual Reality: How New Technologies Can Boost Research on Crowd Dynamics

Mehdi Moussaïd[1]*, Victor R. Schinazi[2], Mubbasir Kapadia[3] and Tyler Thrash[2,4,5]

[1] Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany, [2] Chair of Cognitive Science, Department of Humanities, Social, and Political Sciences, ETH Zurich, Zurich, Switzerland, [3] Computer Science, Rutgers University, The State University of New Jersey, New Brunswick, NJ, United States, [4] Geographic Information Visualization and Analysis, Department of Geography, University of Zurich, Zurich, Switzerland, [5] Digital Society Initiative, University of Zurich, Zurich, Switzerland

The collective behavior of human crowds often exhibits surprisingly regular patterns of movement. These patterns stem from social interactions between pedestrians such as when individuals imitate others, follow their neighbors, avoid collisions with other pedestrians, or push each other. While some of these patterns are beneficial and promote efficient collective motion, others can seriously disrupt the flow, ultimately leading to deadly crowd disasters. Understanding the dynamics of crowd movements can help urban planners manage crowd safety in dense urban areas and develop an understanding of dynamic social systems. However, the study of crowd behavior has been hindered by technical and methodological challenges. Laboratory experiments involving large crowds can be difficult to organize, and quantitative field data collected from surveillance cameras are difficult to evaluate. Nevertheless, crowd research has undergone important developments in the past few years that have led to numerous research opportunities. For example, the development of crowd monitoring based on the virtual signals emitted by pedestrians' smartphones has changed the way researchers collect and analyze live field data. In addition, the use of virtual reality, and multi-user platforms in particular, have paved the way for new types of experiments. In this review, we describe these methodological developments in detail and discuss how these novel technologies can be used to deepen our understanding of crowd behavior.

Keywords: pedestrians, collective movement, complex systems, social interactions, tracking, virtual environment

INTRODUCTION

Understanding crowd movementsis key to the management of dense pedestrian flows in urban areas. Research on crowd dynamics can inform urban planners and help authorities design efficient public places in order to avoid congestions and enhance traffic efficiency (Cassol et al., 2017; Haworth et al., 2017). In addition, crowd research can save lives in extreme situations (Helbing et al., 2014). Recent studies have shown that the frequency and severity of deadly crowd accidents have increased over the past decades (Helbing et al., 2007, 2014; Helbing and Mukerji, 2012). In September 2015, one of the most dramatic crowd stampedes occurred in Mecca during which thousands of pilgrims were crushed to death in a dense crowd (Khan and Noji, 2016). This tragedy is one example of a series of accidents that have occurred in the past decade, costing many lives and undermining trust in public institutions. In the present article, we will describe new technologies that can potentially transform the way crowd researchers address these fundamental issues.

How the System Works

Pedestrian crowds belong to a large family of self-organized social systems (Helbing et al., 2005; Moussaïd et al., 2009), including animal swarms (Camazine, 2003) and human activities such as judgment formation and consumer behaviors (Castellano et al., 2009; Moussaïd et al., 2015). In such systems, the collective dynamics of the group is driven by behavioral propagation processes that are induced by interactions between individuals (Moussaïd et al., 2017). Indeed, pedestrian behaviors tend to spread from person to person, resulting in large-scale snowball effects. For example, when pedestrians slow down or stop in the middle of a dense crowd, they force followers to also slow down or stop in order to avoid a collision. This can trigger a chain reaction as others adapt their movement and/or speed. Behaviors as diverse as choosing an exit door, avoiding others on a particular side, pushing, or escaping from danger are subject to behavioral propagation. This propagation process eventually gives rise to collective patterns, such as lane formation, the emergence of trail networks, and biases in exit choice (Helbing et al., 2005). For example, crowd turbulence is a deadly collective phenomenon that has been recently identified from video surveillance analyses and systematically associated with crowd accidents (Helbing et al., 2007). This pattern is characterized by the occurrence of waves of pushing that propagate from person to person through the crowd. At very high densities, body contacts between neighboring individuals support the spread of pushing forces. These pushing waves set up, merge, and amplify when a certain density threshold is achieved. As a result, people can be trampled by others or crushed against walls. Thus, a large-scale global pattern (e.g., crowd turbulence) can emerge from a simple propagative individual behaviors (e.g., pushing behaviors).

The link between global patterns and the individual behaviors that cause them is often difficult to establish. A crowd is more than a collection of many isolated individuals. Studying individual behaviors in isolation is not sufficient for understanding collective dynamics, and macroscopic descriptions of these patterns are not informative regarding the mechanisms underlying their emergence. Instead, one needs to focus on the causal mechanisms underlying these two levels of observation (i.e., individual and collective behaviors).

How to Study the Crowd

In order to study crowd behavior, researchers use a combination of computer simulations, field observations, and laboratory experiments. Computer simulations explore the conditions in which collective behaviors can emerge by simulating the movements and interactions of many individuals. The outcomes of simulations are determined by behavioral models that describe how individuals respond to their physical and social environments. Existing microscopic pedestrian models include behavioral elements such as how individuals walk to their destinations, how they avoid obstacles, and how they adapt to the presence of other individuals. A large variety of models have been developed in the past. These models include physics-based models (Helbing and Molnár, 1995), biomechanically-based approaches (Singh et al., 2011b), vision-based models (Ondrej et al., 2010; Moussaïd et al., 2011; Dutra et al., 2017),

velocity-based approaches (Guy et al., 2009; van den Berg et al., 2011), and hybrid approaches (Singh et al., 2011a). In addition, macroscopic models aim at describing crowd movement by means of locally averaged quantities, such as the velocity, density, or flow of individuals. This type of model is often inspired by Henderson's original specification with respect to fluid dynamics (Henderson, 1974). The state-of-the-art for crowd modeling techniques has been reviewed in several articles (e.g., Bellomo and Dogbe, 2011; Schadschneider et al., 2011; Degond et al., 2013) and is beyond the scope of this article. A key challenge is to capture the essence of real human crowd behavior while generalizing to future scenarios (e.g., a change in environmental conditions or stress induction in a crowd).

Another methodological approach consists of collecting real-world data directly in the field (e.g., Gallup et al., 2012; Alnabulsi and Drury, 2014). These empirical observations can be used to build data-driven computational models of human crowds (Qiao et al., 2017). Researchers typically set up video recording installations directed at crowded urban environments or use existing recordings from video surveillance platforms. The recorded walking behaviors of pedestrians can then be quantified by reconstructing the positions of individuals from the video images. The advantages of studying real-world phenomena are often undermined by difficulties with the accuracy of these reconstructions, particularly for dense crowds. This quantification step is usually undertaken by means of computer vision software (e.g., Pérez-Escudero et al., 2014) but often requires the tedious efforts of research assistants.

The third approach to studying crowd behavior is to conduct controlled laboratory experiments. In a typical experiment, researchers will invite a group of participants to the laboratory and provide them with specific walking instructions. In the past two decades, a large number of experiments have involved up to hundreds of participants simultaneously, covering a wide range of scenarios. These experiments investigated the study of crowd evacuations, density effects, patterns characterizing uni- and bi-directional flows of people, and large-scale evacuations from public buildings (Hoogendoorn and Daamen, 2005; Jelić et al., 2012; Moussaïd et al., 2012; Burghardt et al., 2013; Wagoum et al., 2017). The popularity of crowd experiments can be explained by the potential to vary experimental factors in a controlled manner, coupled with the ease of tracking participants positions with dedicated tracking devices.

New Perspectives

New technologies such as virtual sensing and multi-user virtual reality platforms can complement the opportunities afforded by field observations and laboratory experiments. Virtual sensing consists of estimating crowd movements by tracking the Wi-Fi and Bluetooth signals emitted by pedestrians' smartphones. Whereas, the idea of estimating a quantity by means of a proxy measure is typically found in other domains (e.g., computer science, chemistry, or transportation science; Liu et al., 2009), this methodology also constitutes a promising line of research for crowd monitoring. In addition, the emergence of multi-user virtual reality platforms can be used to study the movement behavior of crowds instead of individual participants. Controlled

crowd experiments have recently been conducted in virtual environments, extending the limits of possible experimental designs (Thrash et al., 2015; Moussaïd et al., 2016).

We describe how the emergence of virtual sensing and virtual reality can boost crowd research, their potential applications, and corresponding challenges. In the following section, we present previous crowd monitoring techniques and the potential of smartphone-based signals. This section is followed by a discussion of virtual reality from single-user experiments to recent development in multi-user virtual environments. The article concludes with a discussion that highlights the future promises of these techniques for field observations and controlled experiments.

CROWD MONITORING IN THE FIELD

Crowd monitoring involves collecting quantitative information about an existing crowd located in an area of interest, such as crowded streets, music festivals, or train stations. Unlike laboratory experiments and computer simulations, crowd monitoring provides data on real-world behaviors with high external validity. The obtained data may include (i) macroscopic features of the crowd (e.g., density, flow, movement patterns) and/or (ii) microscopic information regarding the pedestrians (e.g., their positions in space, walking trajectories, walking speeds). However, accurate monitoring can be challenging in practice. Crowd monitoring often requires tedious manual corrections and tailored adjustments to specific external factors (e.g., calibrating video analyses techniques to ambient light conditions). There are at least two categories of technical options for monitoring crowds (i.e., conventional methods and virtual sensing).

Conventional Methods

Conventional methods of crowd monitoring include manual crowd counting and computer vision. An early procedure for manual crowd counting was introduced by Herbert Jacobs in 1967—a journalism lecturer at the University of California at Berkeley (Jacobs, 1967). During the Berkeley riots against the Vietnam war, Jacobs observed a crowd from his office window and devised what is known as the "Jacobs method" for estimating its size. The Jacobs method involves estimating the number of people within a square of a stone pavement grid and counting how many of these squares were occupied. Crowd density can then be estimated by calculating the number of people per square meter. This method is still frequently used to estimate crowd density based on video surveillance footage (Raybould et al., 2000). To date, the Jacobs method also remains a simple procedure for extracting the ground truth values used as benchmarks in the validation of more sophisticated methods. Other manual counting approaches include counting people with digital clickers at entrance or exit gates (Bauer et al., 2009, 2011).

Given recent advancements in technology, computer vision techniques have become increasingly popular. This technique consists of extracting relevant crowd information based on the automated analyses of videos. These videos are often sourced from surveillance cameras or aerial images. There are two distinct approaches to computer vision, including the direct approach of detecting people's bodies (Rittscher et al., 2005) or faces (Lin et al., 2001) and the indirect approach of inferring the presence of people using image transformation procedures. For example, researchers have used indirect methods by counting foreground pixels after subtracting the background image (Davies et al., 1995; Ma et al., 2004). Other researchers have employed texture features analysis (Marana et al., 2005), histograms of edge orientations (Dalal and Triggs, 2005), and moving corner points to estimate the number of moving people (Albiol et al., 2009). Crowd flow may also be estimated using the frame difference algorithm (Liang et al., 2014) or the optical flow approach (Andrade et al., 2006).

In the recent years, computer vision techniques have been reshaped by the rise of deep learning (Ouyang and Wang, 2013). Convolutional neural networks can be trained on large hand-annotated crowd datasets (e.g., ImageNet, WWW crowd dataset) to associate image features with higher-level information about the crowd. These methods can produce microscopic quantities, such as the position, number, and trajectories of the pedestrians (Ouyang and Wang, 2012, 2013; Sermanet et al., 2013), or macroscopic information, such as density maps (Sindagi and Patel, 2017), the spatial distribution of the crowd (Kang and Wang, 2014), and contextual information regarding what kind of crowd is present, where the scene occurs, and reasons for the gathering (Shao et al., 2015, 2017). Because deep learning can handle common problems that hinder the efficiency of traditional approaches (e.g., changing camera perspective, body occlusions, and lighting conditions), accuracy levels are typically higher than what can be achieved by conventional methods (Tian et al., 2015).

Despite the fast development of deep learning and the attention it has received in the domain of computer science, this method has not yet widely reached the community of crowd researchers. This is probably related to its lower accessibility for non-experts and the technical complexity of its implementation. To date, traditional crowd monitoring methods remain relatively popular, but the promises of deep learning foreshadow an important development in the near future.

Virtual Sensing

Whereas, conventional methods aim to visually detect the presence of people (with the human eye or the computer eye), virtual sensing consists of detecting traces of people and inferring their numbers, density, and movements. Many methods of crowd sensing rely on emerging technologies that enable the detection of physical and virtual traces left by pedestrians. These methods include carbon dioxide sensors (Ang et al., 2016), audio sensors (Kannan et al., 2012), floor pressure sensors (Mori et al., 2004), seismic sensors (Damarla et al., 2016), motion sensors (Coşkun et al., 2015), and radar sensors (Choi et al., 2016).

In our highly connected world, people do not only leave physical traces in their environment but also emit a variety of virtual traces (e.g., the radio-frequency signals produced by smartphones or other electronic equipments). The increased reliance on smartphones and other connected devices has motivated researchers to extract the crowd information provided by these mobile devices (Eagle et al., 2009; Ding et al.,

2015). Numerous applications have been developed that employ smartphones as sensors for the recognition of activities such as mobility, health information, and social interactions (see survey in Khan et al., 2013). In the specific case of crowd sensing, collecting location data from smartphones can be achieved by accessing a device's GPS or Wi-Fi positioning information (with positional accuracies of ∼5 and 20 m, respectively; Azizyan et al., 2009; Van Diggelen, 2009).

However, collecting position information is not trivial. For privacy reasons, the positioning information of any randomly selected pedestrian is typically not publicly available. Some researchers have circumvented this challenge by setting up a voluntary participatory system. Here, volunteers can register to participate in the study and install an experimental application on their smartphone. The application continuously records the user's spatial position and sends it to a central server. Recent studies have shown that many individuals are willing to install such an application and share this type of data as long as the scientific use of this data is communicated in a transparent manner and when the participants can receive valuable information in return (Wirz et al., 2012). For example, participants recruited at a music festival may be able to use the application to access an interactive program guide, a map of the neighboring points of interest, background information regarding ongoing concerts, and other social features. Furthermore, the application can be used to send personalized location-dependent information to the users. For example, the police can inform attendees located in a particular area about how to behave in case of an emergency. One challenge of this sensing method is that the researcher cannot expect to receive position information from all individuals in the area of interest. Because only a fraction of people will be using the application, researchers must extrapolate the positions and movements of the entire crowd from those of the collected sample.

This method has been previously employed during the 2011 Lord Mayor Show in London (Wirz et al., 2013) and the 2013 Züri Fäscht in Zürich (Blanke et al., 2014). During the Lord Mayor Show, 828 users downloaded the application (out of nearly half a million visitors) and ∼4 millions GPS positions were collected at a sampling rate of 1 Hz. This method was validated by comparing the GPS position data to a ground truth sample resulting from the semi-automatic monitoring of surveillance camera recordings. This study demonstrated that the application users were distributed across the festival area similarly to the rest of the crowd. Indeed, there was a positive correlation between the density of application users and the actual crowd density (**Figure 1A**). As an illustrative result, **Figure 1B** shows a map of the crowd density in the festival area.

For the 2013 Züri Fäscht—a 3-day event comprising concerts and shows—the scaling was considerably increased. Out of 2 million total visitors, 28,000 users downloaded the application, resulting in ∼25 million location updates. The higher participation rate of this second deployment resulted from an important marketing effort in promoting and distributing the application. Several functionalities were added, including a "friend finder" that allowed users to locate their friends in the event they became lost in the crowd. The gamification of this application (with a "trophy collector" function) also possibly contributed to the higher download rate. Finally, a link to the user's Facebook profile favored the viral propagation of the application on social networks.

Overall, this application allowed for the collection of detailed data regarding the crowd at a scale and with an accuracy that was rarely achieved in the past.

Despite the advantages of virtual sensing with active participants, this method relies on an intensive marketing effort. Alternatively, researchers may track pedestrians passively using the Bluetooth and Wi-Fi signals emitted by their mobile devices. Indeed, Bluetooth and Wi-Fi signals can be detected using dedicated scanners (Musa and Eriksson, 2012; Barbera et al., 2013). When applied to crowd observation, stationary scanners positioned in the area of interest can allow the detection of virtual traces left by pedestrians and thus the estimation of their number and displacement (Fukuzaki et al., 2014; Schauer et al., 2014). Hence, pedestrians do not need to actively cooperate with the researchers by downloading an application on their phones.

However, the deployment of the scanners can be challenging. One important issue with Wi-Fi is the interruption of the signal propagation path caused by solid obstacles located between the source and the scanners. In addition, human bodies can also produce a shield effect that causes fluctuations in the signal. One solution is to mount them *above* the crowd, thus enabling a free line of sight toward the devices. While this solution is easily applicable in indoor environments, it is more challenging when tracking people in open spaces, such as commercial walkways or music festivals.

Virtual sensing with passive participants has been successfully deployed several times in the past (e.g., in shopping malls, car exhibitions, and airports, see Fukuzaki et al., 2014; Schauer et al., 2014). For example, Weppner et al. (2016) used a setup consisting of 31 scanners (covering a total area of ∼6,000 m^2) during the IAA car exhibition in Frankfurt. Data was collected for 13 business days, producing nearly 90 million data points from a total of over 300,000 unique mobile devices. A video-based manual counting procedure was also employed in order to validate the virtual sensing data. The scanners were mounted on the ceiling with an average distance of 14 m between them and an average scanning zone of 180 m^2 for each of them.

Whenever, pedestrians walked through the detection area, the Bluetooth and Wi-Fi signals emitted by their mobile devices were detected by the scanners and sent to a central database server. Every incoming signal was associated to an RSSI value (i.e., the Received Signal Strength Indication). This information can be combined with the coordinates of the scanners to estimate the location of the pedestrian during a post-processing phase. Multiple scanners can detect the presence of a unique mobile device at a given moment of time. The simplest localization method is to assign the spatial coordinates of the scanner that has recorded the highest RSSI (i.e., the strongest signal) to the pedestrian. A more sophisticated method is based on an RSSI-weighted average of the scanners locations. In a preliminary accuracy evaluation phase, the positioning error was estimated to a maximum of 10 m for 90% of the devices.
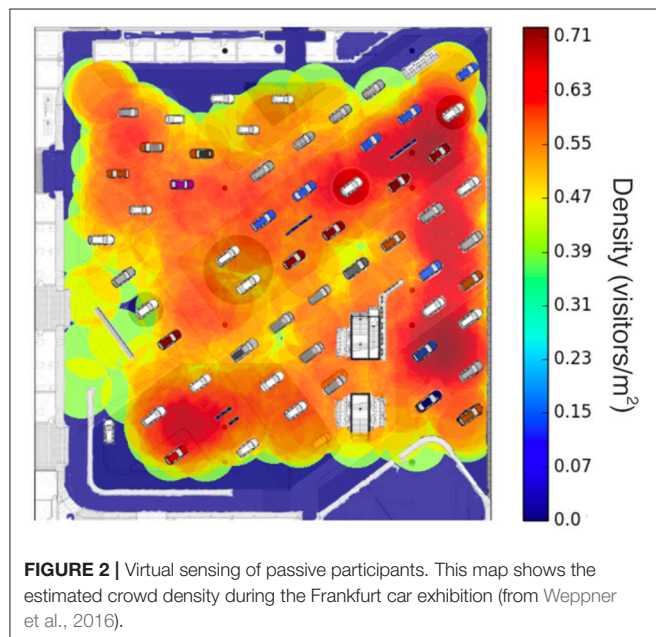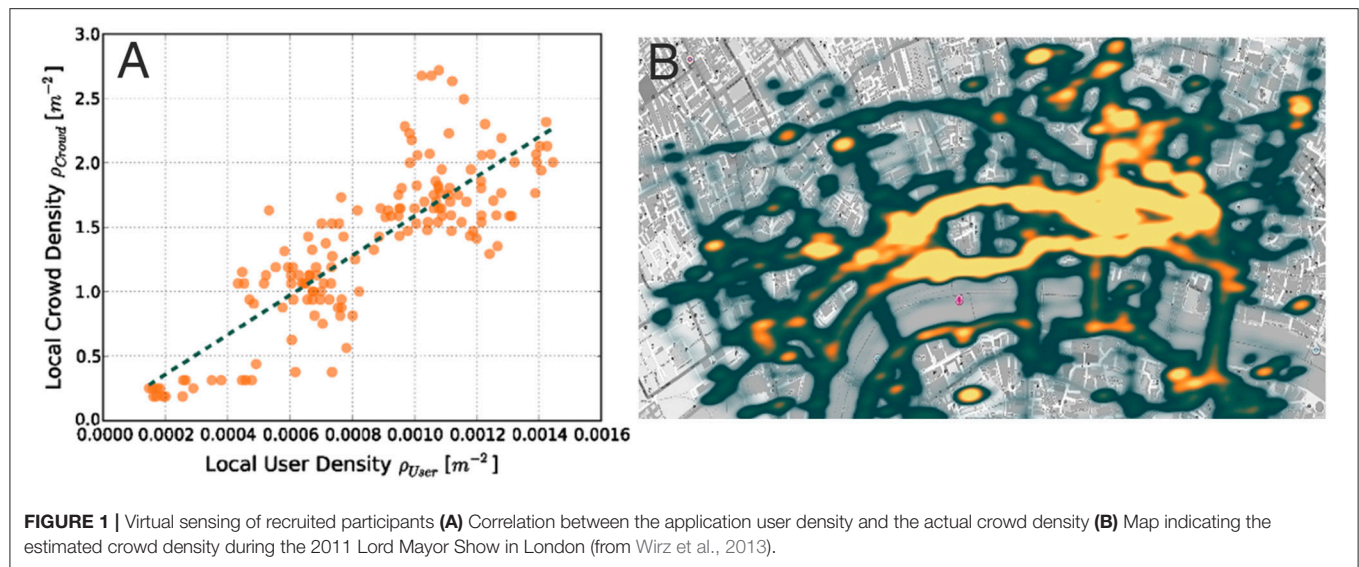
FIGURE 1 | Virtual sensing of recruited participants (A) Correlation between the application user density and the actual crowd density (B) Map indicating the estimated crowd density during the 2011 Lord Mayor Show in London (from Wirz et al., 2013).

FIGURE 2 | Virtual sensing of passive participants. This map shows the estimated crowd density during the Frankfurt car exhibition (from Weppner et al., 2016).

Figure 2 shows the estimation of local densities in subregions delimited by the boundaries of a Voronoi cell surrounding each scanner.

Calibration was necessary to convert the estimated density of people into the actual density because not all visitors were carrying a detectable device and the signal was not always detected. Toward this end, ground truth manual measurements were compared to the measures provided by the sensors. Weppner et al. (2016) calculated that the measures provided by the sensors have to be multiplied by an average of 1.5 in order to match the ground truth values. In practice, the value of the multiplier might vary depending on social and

environmental conditions and would need to be calibrated by means of preliminary evaluation data.

VIRTUAL REALITY IN THE LABORATORY

Virtual reality (VR) is a technology that involves presenting a person with a responsive artificial environment. Participants in VR studies are typically able to look around, move in, and interact with the virtual environment. As such, VR constitutes an interesting opportunity to study pedestrians' behaviors such as locomotion (i.e., bodily movement through the immediate environment) and wayfinding (i.e., spatial decision-making in a large-scale environment; Montello, 2005).

Techniques and Single-User Experiments

In VR, the interaction between a navigator and the environment is mediated by a display (e.g., projection screen, head-mounted display), and a control interface (e.g., a joystick, a mouse, and keyboard, head movement sensors). Large projection screens and desktop displays often provide a more natural field of view but do not always allow users to rotate their bodies 360° in order to experience the virtual environment (but see Höllerer et al., 2007). In contrast, head-mounted displays (HMDs) are relatively mobile and restrict visual access to the external world (e.g., Oculus Rift, https://www.oculus.com/; HTC Vive, https://www.vive.com/us/) (see e.g., Chance et al., 1998; Waller et al., 2004; Foo et al., 2005; Kinateder and Warren, 2016). One consequence of using VR displays is that distances are systematically underestimated to a greater extent than distances estimated in the real world (Knapp, 2003). However, training in VR that involves explicit visual feedback can reduce these biases (Richardson and Waller, 2005). Similarly, spatial updating has been found to be less precise in VR without physical turns (Klatzky et al., 1998), but biases in turn perception *per se* can be reduced with explicit visual feedback (Bakker et al., 2001).

The control interface translates the movements of users into visual feedback on the display. Two important aspects of control interfaces are the position of the body (Taube et al., 2013) and the possible ways in which specific actions (e.g., pushing a joystick forward) are connected with specific types of feedback (e.g., the expansion of optic flow). During locomotion in VR, the user's body can be sitting (e.g., Richardson et al., 1999), lying (as in neuroscientific research; Taube et al., 2013), or standing (e.g., Nescher et al., 2014). While sitting or lying (or standing in place), the user does not receive proprioceptive (i.e., body-based) feedback. In addition, lying causes a conflict in perspective between facing upwards in the real environment (e.g., the fMRI scanner) and facing forward in the virtual environment (Taube et al., 2013). Comparisons of control interfaces are often case-specific. For example, Thrash et al. (2015) found that users' performance on navigation-related tasks was more efficient and less error prone with a mouse-and-keyboard setup than a handheld joystick. However, less attention has been allocated to theoretical explanations for why users tend to perform better with some interfaces than others. While mouse-and-keyboard setups are often more familiar than joysticks, the extent to which one interface is more "intuitive" than the other is unknown (Lapointe et al., 2011). This challenge may be addressed in the future by studies that focus on the impact of training on interface use or on how to allow for realistic walking in VR.

For realistic walking, some researchers have employed omnidirectional treadmills (as a hardware solution; e.g., Souman et al., 2010) and redirected walking algorithms (as a software solution; e.g., Razzaque et al., 2001). Redirected walking steers users toward particular targets by expanding and compressing rotations and translations and allows for locomotion through environments that are larger than the external infrastructure. Even when VR participants walk with an HMD (without these visual distortions), the HMD necessarily translates head movements into visual feedback and thus constitutes a control interface.

Advancements in control interface technology will be critical for studies of locomotion but may be less critical for studying certain aspects of wayfinding. Indeed, during wayfinding, the decisions executed by the navigator typically depend less on physical collisions or maneuverability than incomplete mental representations and salient environmental cues.

Wayfinding behavior can be classified as either path integration or landmark-based (Taube et al., 2013). During path integration, observers rely on idiothetic cues in order to maintain their orientations and positions during movement through a large-scale environment (Gallistel, 1990). Landmark-based navigation relies primarily on allothetic cues (e.g., visible objects along a route; Presson and Montello, 1988) and is associated with scene processing (Epstein and Vass, 2014) and survey representation (Kitchin and Blades, 2002). Indeed, this type of wayfinding has been successfully studied using a variety of VR systems, including projection screens in fMRI scanners (Epstein et al., 2017), desktop displays with simple controls (Waller and Lippa, 2007), and HMDs with naturalistic walking (Hodgson et al., 2011).

Virtual reality has allowed real humans to interact with their digital counterparts (i.e., avatars) in an effort to study more detailed local interactions under controlled experimental conditions. For example, Olivier and colleagues have used VR in order to study how people avoid collisions with groups (Bruneau et al., 2015), the impact of social roles on collision-avoidance strategies (Olivier et al., 2013), as well as human-robot interactions (Vassallo et al., 2017). Similarly, Warren and colleagues have focused on human locomotion and spatial navigation using VR (Bonneaud et al., 2012). These studies have allowed researchers to test theories of perceptual-motor control and develop a formal model of pedestrian behavior (Warren and Fajen, 2004; Bonneaud and Warren, 2012). This model has been expanded to include perception (Bruggeman et al., 2007; Warren and Fajen, 2008) and behaviors such as target interception (Fajen and Warren, 2007) and collision avoidance with both static and moving objects (Fink et al., 2007).

Immersive Multi-User Experiments

One drawback of single-user experiments is the lack of interactions between participants. The collective dynamics of a crowd cannot be explained by the accumulation of many isolated individuals. Rather, collective behaviors stem from social interactions between pedestrians. Observing the interactions of a single participant with simulated agents constitutes an interesting step toward studying crowd dynamics in VR (Drury et al., 2009). Nevertheless, insight into collective behavior remains elusive because the dynamics of the group are largely determined by the behavior of the virtual agents implemented by the experimenter.

This challenge has been recently addressed with the development of *multi-user* virtual environments (Normoyle et al., 2012; Bode and Codling, 2013; Bode et al., 2014; Carlson et al., 2014; Moussaïd et al., 2016; Boos et al., 2017). These multi-user environments enable the observation of a *crowd* of participants moving and interacting in a shared virtual environment simultaneously. In a typical multi-user experiment, every participant controls an avatar in the virtual environment from a first-person perspective. The avatars can view and interact with each other in real time (e.g., avoiding, following, or colliding) and thus mimic some aspects of social interactions among real pedestrians. In the following series of experiments, Moussaïd et al. (2016) explored the potential of multi-user VR using desktop displays with a mouse-and-keyboard control interface (Thrash et al., 2015).

Validation

Given the novelty of multi-user VR experiments, initial research has focused on validating simple crowd behaviors observed in virtual worlds. Here, we describe two studies that have compared avoidance maneuvers and simple evacuation situations against real-world data.

Side Preference

Avoidance maneuvers between pedestrians are characterized by a well-known social bias called the side preference (Helbing, 1992). In most Western countries, people preferentially evade each other on the right-hand side. This bias is a *social* attribute
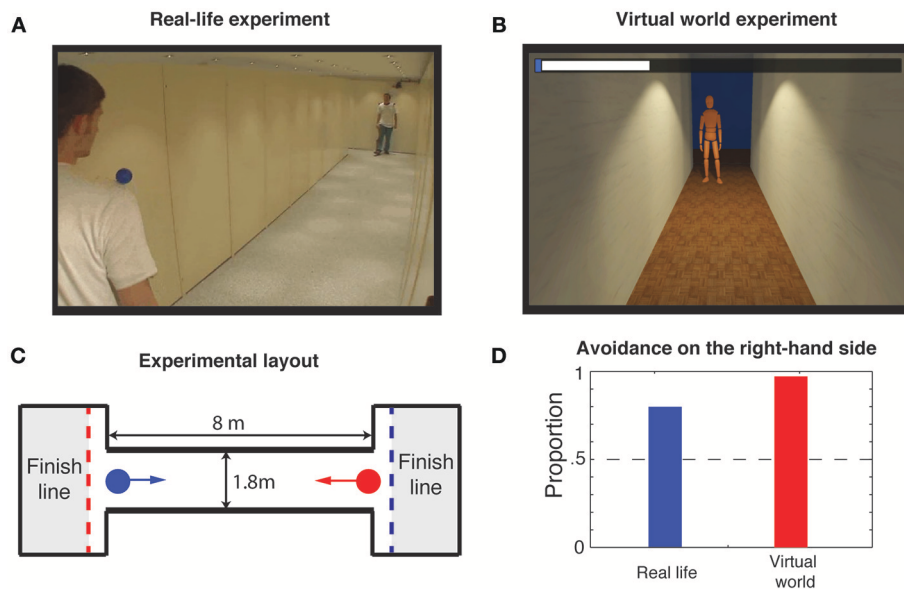
FIGURE 3 | The side preference. **(A)** Illustration of the real-life experiment in which pairs of participants were instructed to avoid each other in a narrow corridor (Moussaïd et al., 2009). **(B)** Replication of the side preference experiment in the multi-user virtual environment. **(C)** The layout and the dimensions of the corridor are identical between the two experiments. **(D)** In both experiments, participants exhibited a marked preference for avoidance on the right-hand side (81% in the real-life, 95% in the virtual environment), demonstrating that people rely on the same social norms in real and virtual settings.

that does not occur during the avoidance of a static obstacle (Moussaïd et al., 2009). In a multi-user VR experiment, 95% of the participants exhibited the side preference, compared to 81% in an identical real-world study (Moussaïd et al., 2009, 2016; **Figure 3**). This suggest that participants in VR can consider other avatars as "real" people and expect them to follow similar social norms.

Simple Evacuation

The second validation experiment focused on evacuation dynamics. Previous research has demonstrated that the outflow during an evacuation of a group of people increased linearly with the width of the room doorway (Kretz et al., 2006; Liddle et al., 2009; Seyfried et al., 2009; Daamen and Hoogendoorn, 2010).
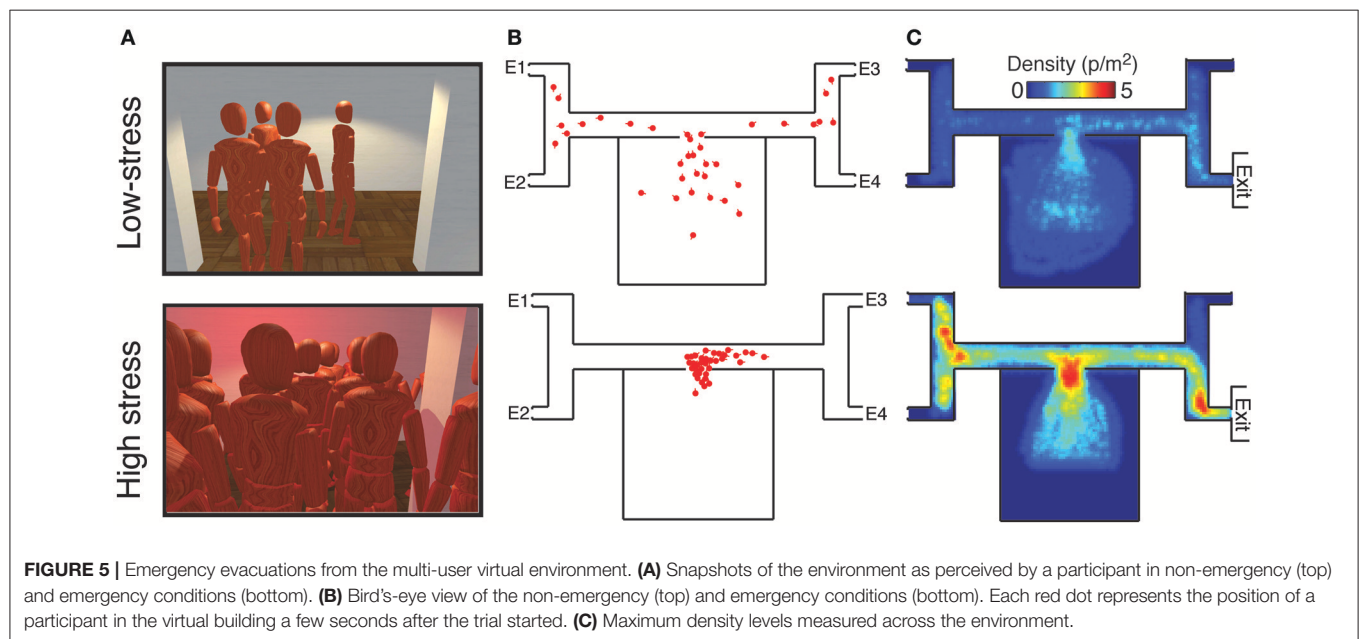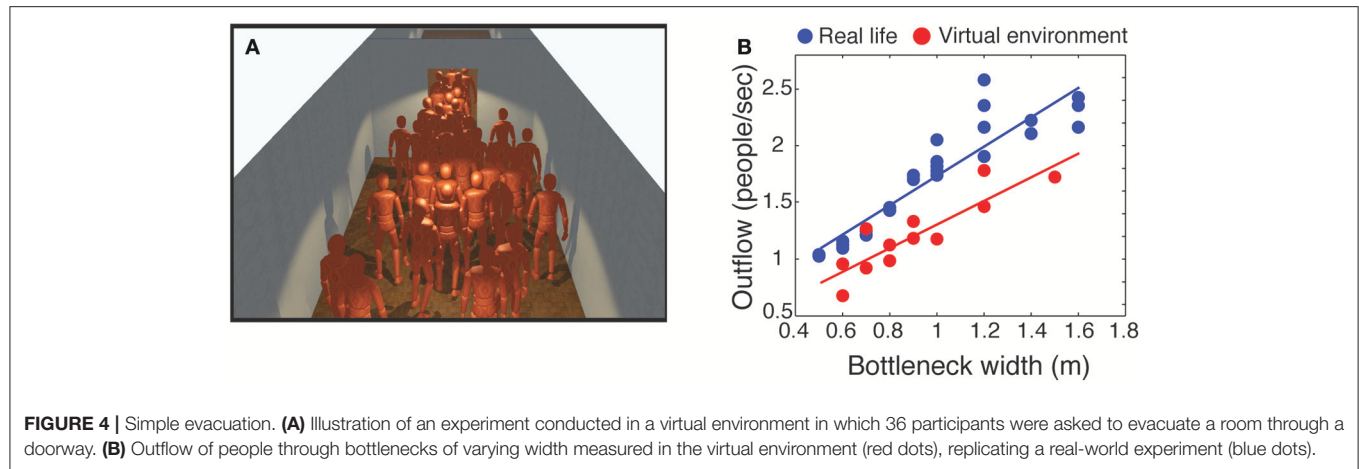
One of these evacuation experiments has been replicated in desktop VR (Moussaïd et al., 2016). A total of 36 participants were immersed simultaneously in a large virtual room and instructed to evacuate through a doorway of varying width (Kretz et al., 2006). Consistent with real-world findings, the outflow of pedestrians increased linearly with the bottleneck width (**Figure 4**). However, compared to a larger body of real-world datasets, the outflow of participants was smaller in the virtual environment. This difference can be attributed to micro-navigation factors such as differences in walking speed, acceleration, and/or shoulder movements.

Emergency Evacuations

Multi-user virtual environments also offer the advantage of enabling the investigation of difficult (if not impossible) scenarios. For example, the collective behavior that occurs during emergency situations (e.g., evacuating a burning building) can

be challenging to study in the real world because of ethical and safety reasons (Schadschneider et al., 2011). Recently, emergency evacuations were investigated in virtual settings. Large groups of participants were instructed to evacuate a virtual building with four possible exits, only one which was not blocked by fire. For each trial, the location of the correct exit was randomly chosen, and only a randomly selected subset of participants were told which exit was correct. We compared collective behaviors between non-emergency and stressful emergency conditions. In the study, the two conditions differed by three factors. Specifically, in the stressful emergency condition, there was a short time limit imposed, participants were penalized for not finding the correct exit, and the environment contained stressful elements such as red blinking lights and a siren. In contrast, in the non-emergency condition, no time limit was imposed, participants were rewarded for finding the correct exit, and the environment lacked blinking lights and sirens.

The results revealed significant differences between the two conditions (**Figure 5**). While participants searched for the exit in a slow and orderly manner in the non-emergency condition, mass herdings and severe crowding occurred in the emergency condition. In particular, in the non-emergency condition, participants tended to stay reasonably safe distances from one another in order to avoid a monetary penalty for colliding with each other. In contrast, a high number of collisions occurred in the high-stress condition, despite having the same collision penalty. Density levels remained lower than 2 people per m^2 in the non-emergency condition, as typically observed in everyday congested zones (Still, 2000). Under high stress, the density level reached values up to 5 people per m^2. This value is close to the critical threshold of crowd turbulence, a deadly collective

FIGURE 4 | Simple evacuation. **(A)** Illustration of an experiment conducted in a virtual environment in which 36 participants were asked to evacuate a room through a doorway. **(B)** Outflow of people through bottlenecks of varying width measured in the virtual environment (red dots), replicating a real-world experiment (blue dots).

FIGURE 5 | Emergency evacuations from the multi-user virtual environment. **(A)** Snapshots of the environment as perceived by a participant in non-emergency (top) and emergency conditions (bottom). **(B)** Bird's-eye view of the non-emergency (top) and emergency conditions (bottom). Each red dot represents the position of a participant in the virtual building a few seconds after the trial started. **(C)** Maximum density levels measured across the environment.

phenomenon (Helbing et al., 2007). Another collective pattern that emerged in the emergency condition was herding. While participants in the non-emergency condition tended to choose a random branch at each intersection, the majority of participants herded in the same direction in the emergency condition, which amplified the crowding pattern.

The development of multi-user virtual environments for conducting crowd experiments is promising but still at the early stages. Additional validation experiments should be conducted. In addition, there are necessary improvements with respect to simulating social and physical interactions between avatars during navigation. For example, social interactions may be impacted by appearance and behavioral realism of other avatars in the virtual environment (e.g., gait; Narang et al., 2017b). These aspects of realism in VR can be improved using new methods for generating avatar movement based on the recordings of real people (Narang et al., 2017b). Empirical research has also

demonstrated that the match between appearance and behavioral realism is critical for recognizing one's own movement (Narang et al., 2017a) and co-presence (Bailenson et al., 2005). With respect to crowds, Prazak and O'Sullivan (Pražák and O'Sullivan, 2011) suggest that the crowd's perceived realism depends on the number of animations particular to individual avatars.

Additional challenges for multi-user VR include the lack of haptic feedback and sound rendering, material constraints associated with equipping multiple participants with individual displays (e.g., HMDs) and controls, and sufficient training with these controls. Previous research has suggested the benefits (e.g., improved immersion) of haptic feedback using haptic garments (Ryu and Kim, 2004), vibrating actuators (Louison et al., 2017), and quadcopters (among others; see Knierim et al., 1998). Similarly, the rendering of spatialized sounds may complement visual feedback by providing temporal information that can improve presence (see Serafin et al., 2015 for a review). However,

both haptic feedback and sound rendering require additional computing power and impose material constraints. For example, equipping 36 participants with HMDs, haptic garments, and spatialized sounds would be prohibitively expensive in terms of finances and computational resources. These constraints also require participants in multi-user VR to use simple controls such as a joystick. Training with these controls is critical given that participants may need to negotiate both static (e.g., walls) and dynamic (e.g., other avatars) obstacles (Grübel et al., 2017).

Alternative Approaches to Multi-User Experiments

Other approaches have been used to study crowd behaviors in VR. Compared to the above examples, these approaches are not presented from a first-person perspective and implement a less realistic graphical environment.

One of the first attempts to study evacuations in multi-user virtual environments was conducted within the popular massive multiplayer online game "Second Life." There, users can create an avatar, explore a large virtual environment, and interact with other users' avatars. Whereas the primary purpose of Second Life is entertainment, researchers have used it to conduct behavioral experiments (Molka-Danielsen and Chabada, 2010; Normoyle et al., 2012). For these experiments, participants were recruited among existing users of Second Life with announcements posted in the virtual world. Participants met in a virtual building and then were asked to evacuate because of a virtual fire. The experimenters were able to characterize numerous aspects of emergency evacuations (e.g., exit choice, knowledge about the building plan), but this type of experimental setup offers little experimental control. Nevertheless, using an existing virtual world already populated with thousands of users could potentially allow the development of very large-scale experiments (i.e., with more than 36 participants). In addition, other massive multiplayer online platforms may allow for a combination of both larger crowds and experimental control to study phenomena such as crowd disasters.

Other simpler approaches for conducting crowd experiments in virtual environments have also been developed. For example, Bode and Codling have studied various aspects of evacuation dynamics by having participants control the movement of a dot with a computer mouse through a two-dimensional environment from a top-down perspective (Bode and Codling, 2013; Bode et al., 2014, 2015). The authors managed to highlight some important aspects of participant behavior during evacuations, such as the impact of congestions, static signs, social cues, and memorized information on routing and exit choice dynamics. Although these experiments were designed for a single participant interacting with simulated agents, adapting this approach to multiple simultaneous users should only present minor technical challenges.

Similarly, the HoneyComb paradigm has a multi-player design in which each participant controls a dot on a two-dimensional playfield (Boos et al., 2017). Using their mouse, groups of participants can navigate simultaneously in a shared environment. Every individual can see the position and the

movement of the those who are located within a particular perceptual radius. In such a way, researchers investigated a series of fundamental questions related to the role of leadership (Boos et al., 2014), spatial attraction (Belz et al., 2013), and competition (Boos et al., 2015) on collective flocking patterns.

DISCUSSION

Conventional methods of crowd monitoring are difficult to implement for tracking large crowds, and experimental approaches often face organizational and ethical challenges. Owing to recent technological developments, novel methods of crowd monitoring (i.e., virtual sensing) and crowd experimentation (i.e., multi-user virtual reality) have emerged and constitute promising complementary options for crowd researchers.

Virtual Sensing

Most pedestrians carry a connected device (e.g., a smartphone) that continuously emits radio-frequency signals. Whereas, the physical locations of individuals are often difficult to establish using video recordings, these locations can be inferred by detecting and tracking the virtual traces left by their devices. Crowd monitoring techniques have rapidly evolved from manual counting to computer-based video analyses. Researchers can now transition toward virtual sensing techniques. However, two major challenges for this approach are to access a sufficiently large proportion of these signals and to estimate their locations as accurately as possible.

Toward this end, two methods have been developed. The first method consists of distributing a dedicated application to a large sample of users. This application can continuously record users' positions and send these positions to a central server. The second method consists of monitoring the Wi-Fi signals emitted by devices using dedicated sensors installed in the area of interest. Both methods are able to accurately represent the crowd's movement and density. However, both methods also require a considerable amount of effort to set up. Deploying an application requires a marketing effort to distribute as broadly as possible and convince people to install and activate it. Remarkable progress has been made in that regard between the two past deployments (at the Lord Mayor Show in London and the 2013 Züri Fäscht in Zürich) of a virtual sensing system, for which the number of participants has increased from 828 users to 28,000 users. In particular, the authors of these studies noticed that the application should offer a variety of services to the users, explicitly communicate about what usage is made with the collected data, and make use of social networks and social recommandation tools.

In contrast, monitoring Wi-Fi signals does not require the explicit cooperation of the individuals. However, dedicated signal sensors must be installed in the area of interest and may require permission from the event organizers. In addition, the sensors must be positioned as much above the crowd as possible in order to avoid signal interruptions and obstructions. Recently, innovations in animal tracking have demonstrated the advantages of using drones to collect video and GPS data

on the movement of wild baboons (Strandburg-Peshkin et al., 2017). Similarly, one could imagine embedding radio-frequency sensors in drones flying above the crowd, which could minimize signal interruptions and convert the sensor into a mobile installation.

Another branch of virtual sensing employs the traces left by interactions between people on the phone or the Internet. To date, such an approach has been used to collect macroscopic data such as unemployment levels, disease prevalence, and consumers behavior based on Internet search queries (Ginsberg et al., 2008; Goel et al., 2010). People's positions in space can also be inferred from their activity patterns. For example, Gonzalez and colleagues used the data from a mobile phone carrier containing the date, time, and coordinates of the phone towers routing the phone calls of ∼6 million users (González et al., 2008). The movements of each user was then inferred by tracking the locations of the phone towers routing the communications despite low spatial resolution (∼3 km^2) and restrictions regarding data accessibility. Nevertheless, it has been shown that the spatial density of phone communications correlated with the volume of geolocalized tweets recorded over the same period on Twitter (Botta et al., 2015). In other words, the number of tweets and the place where they were produced—free and easily accessible data—can serve as a proxy to estimate the density of people in a certain area of interest.

In general, virtual sensing approaches remain less accurate than conventional video-based tracking methods. The positioning of the individuals is, at best, estimated within a few meters of uncertainty. This challenges the extraction of individual-level mechanisms underlying the crowd dynamics. However, virtual sensing has a larger spatial and temporal reach, potentially covering an entire city during unlimited time periods. As such, both methods complement each other well and should eventually constitute different options in the crowd researcher's toolbox.

Virtual Reality

While virtual sensing allows for the observation of natural crowds, multi-user virtual reality provides more control over experimental conditions and the ability to draw causal inferences. This approach builds on single-user virtual reality by allowing for the study of simultaneously immersed users. These multi-user virtual environments have several other advantages.

First, virtual environments are easy to manipulate. Researchers can conduct experiments in virtual buildings, streets, stadiums, or large vehicles such as planes and boats of different typologies and sizes. Unlike real-world experiments that rely on existing physical infrastructures, virtual designs can modify existing environments or create new ones. For example, the side preference experiment described above was conducted both in the real world and in the virtual environment. In the real world, 144 replications of the experiment were collected during several days. In the virtual world, 561 replications of the same experiment were collected in <15 min. Regarding the creation of new environments or situations, experiments

can be conducted to address questions that were previously unapproachable because of safety or ethical issues. For example, they enable the systematic investigation of crowd behavior under stressful and dangerous conditions with real human participants.

Second, multi-user virtual environments allow for greater experimental control. For example, experimental variables such as light level, walking speed, and body size may be manipulated in a way that is not possible in real-world settings. Experiments could also modify real participant behavior to create artificial agents and induce the propagation of certain behaviors through the crowd.

Third, experiments in multi-user virtual environments allows the collection of a large variety of measurement variables with high precision. Participants' positions, speeds, and body and head orientations can be easily captured at high resolution and with minimal measurement errors. In addition, other types of behavior could also be measured such as properties of participants' gaze (using eye trackers) and their physiological states (using electrocardiograms or skin conductance sensors).

While some researchers have studied crowd behavior online with mixed results, new technologies may allow for carefully controlled multi-user experiments in the near future. In such scenarios, participants could use their own computer setup and participate from home. It may be unrealistic to expect a large number of participants at their respective homes wearing HMDs for research purposes. However, desktop computers with mouse and keyboard setups may be sufficient for some experiments similar to those already conducted in the laboratory. These advancements suggest that massive online crowd experiments could be used for studying thousands of participants connected to an experimental server at a given moment of time. Previously, similar group experiments were conducted in the fields of social psychology and network science (Mason and Watts, 2012; Mao et al., 2016). In these experiments, up to a 100 of online participants were tested simultaneously, This approach could be further facilitated by the existence of crowdsourcing platforms for recruiting participants such as Amazon Mechanical Turk or Prolific Academic (Mason and Suri, 2012).

Despite these advantages, multi-user virtual reality cannot be considered as a replacement for conventional real-world experiments. It offers some advantages, like a greater control on external variables, the ease of designing environments, and the potential for exploring dangerous situations, but also has drawbacks. For example, the feeling of body contacts in high density situations is difficult to communicate realistically. Similarly, there exist numerous micro-navigation differences that prevent participants from modulating their speed and acceleration as they would in real life.

CONCLUSION AND PERSPECTIVES

In this review, we have described two technological innovations that can offer promising new perspectives for crowd researchers.

While live monitoring techniques can facilitate data collection for field studies, multi-user virtual reality offers new opportunities for conducting experiments with greater flexibility and control. Similar developments are also taking place for the study of other self-organized social systems. Animal tracking methods are currently undergoing major changes with the development of high-accuracy GPS methods (e.g., Nagy et al., 2010; Strandburg-Peshkin et al., 2015, 2017). At the same time, virtual reality is emerging as a powerful tool for studying social interactions among fish and understanding the resulting collective behaviors of the school (Ioannou et al., 2012; Stowers et al., 2017). The parallel development of virtual sensing and virtual reality across different social systems confirms the important role that these two methods might play for the study of the self-organized crowd phenomena in the future.

AUTHOR CONTRIBUTIONS

MM conceived and structured the review. MM, VS, MK, and TT wrote the paper.

ACKNOWLEDGMENTS

REFERENCES

Albiol, A., Silla, M. J., Albiol, A., Mossi, J. M. (2009). "Video analysis using corner motion statistics," in *IEEE International Workshop on Performance Evaluation of Tracking, and Surveillance,* 31–38.

Alnabulsi, H., and Drury, J. (2014). Social identification moderates the effect of crowd density on safety at the Hajj. *Proc. Natl. Acad. Sci. U.S.A.* 111, 9091–9096. doi: 10.1073/pnas.1404953111

Andrade, E. L., Blunsden, S., and Fisher, R. B. (2006). "Hidden markov models for optical flow analysis in crowds," in *18th International Conference on Pattern Recognition (ICPR'06),* 460–463.

Ang, I. B. A., Salim, F. D., and Hamilton, M. (2016). "Human occupancy recognition with multivariate ambient sensors," *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops),* 1–6.

Azizyan, M., Constandache, I., and Roy Choudhury, R. (2009). "Surroundsense: mobile phone localization via ambience fingerprinting," in *Proceedings of the 15th Annual International Conference on Mobile Computing and Networking* (New York, NY: ACM), 261–272.

Bailenson, J. N., Swinth, K., Hoyt, C., Persky, S., Dimov, A., and Blascovich, J. (2005). The independent and interactive effects of embodied-agent appearance and behavior on self-report, cognitive, and behavioral markers of copresence in immersive virtual environments. *Presence* 14, 379–393. doi: 10.1162/105474605774785235

Bakker, N. H., Werkhoven, P. J., and Passenier, P. O. (2001). Calibrating visual path integration in ves. presence: teleoperators and virtual environments. *MIT Press* 10, 216–224. doi: 10.1162/105474601750216812

Barbera, M. V., Epasto, A., Mei, A., Perta, V. C., and Stefa, J. (2013). "Signals from the crowd: uncovering social relationships through smartphone probes," in *Proceedings of the* 2013 *Conference on Internet Measurement Conference* (New York, NY: ACM), 265–276.

Bauer, D., Brändle, N., Seer, S., Ray, M., and Kitazawa, K. (2009). "Measurement of pedestrian movements: a comparative study on various existing systems," in *Pedestrian Behaviour: Models, Data Collection and Applications,* ed H. Timmermans (Bingley: Emerald), 301–320.

Bauer, D., Ray, M., and Seer, S. (2011). Simple sensors used for measuring service times and counting pedestrians: strengths and weaknesses. *J. Transport. Res. Board* 2214, 77–84. doi: 10.3141/2214-10

Bellomo, N., and Dogbe, C. (2011). On the modeling of traffic and crowds: a survey of models, speculations, and perspectives. *SIAM Rev.* 53, 409–463. doi: 10.1137/090746677

Belz, M., Pyritz, L. W., and Boos, M. (2013). Spontaneous flocking in human groups. *Behav. Process.* 92, 6–14. doi: 10.1016/j.beproc.2012.09.004

Blanke, U., Tröster, G., Franke, T., and Lukowicz, P. (2014). "Capturing crowd dynamics at large scale events using participatory GPS-localization," in *2014*

IEEE Ninth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 1–7.

Bode, N., and Codling, E. (2013). Human exit route choice in virtual crowd evacuations. *Anim. Behav.* 86, 347–358. doi: 10.1016/j.anbehav.2013.05.025

Bode, N. W., Kemloh Wagoum, A. U., and Codling, E. A. (2014). Human responses to multiple sources of directional information in virtual crowd evacuations. *J. R. Soc. Interface* 11:20130904. doi: 10.1098/rsif.2013.0904

Bode, N. W., Miller, J., O'Gorman, R., and Codling, E. A. (2015). Increased costs reduce reciprocal helping behaviour of humans in a virtual evacuation experiment. *Sci Rep.* 5:15896. doi: 10.1038/srep15896

Bonneaud, S., Rio, K., Chevaillier, P., and Warren, W. H. (2012). "Accounting for patterns of collective behavior in crowd locomotor dynamics for realistic simulations," in *Transactions on Edutainment, VII,* eds Z. Pan, A. D. Cheok, W. Muller, M. Chang, and M. Zhang (Berlin; Heidelberg: Springer), 1–11.

Bonneaud, S., and Warren, W. H. (2012). "A behavioral dynamics approach to modeling realistic pedestrian behavior," in *6th International Conference on Pedestrian and Evacuation Dynamics,* 1–14.

Boos, M., Franiel, X., and Belz, M. (2015). Competition in human groups—Impact on group cohesion, perceived stress and outcome satisfaction. *Behav. Process.* 120, 64–68. doi: 10.1016/j.beproc.2015.07.011

Boos, M., Li, W., and Pritz, J. (2017). "Patterns of group movement on a virtual playfield: empirical and simulation approaches," in *Social Network Analysis: Interdisciplinary Approaches and Case Studies* (Boca Raton, FL: CRC Press), 197.

Boos, M., Pritz, J., Lange, S., and Belz, M. (2014). Leadership in moving human groups. *PLoS Comput Biol.* 10:e1003541. doi: 10.1371/journal.pcbi.1003541

Botta, F., Moat, H. S., and Preis, T. (2015). Quantifying crowd size with mobile phone and Twitter data. *R. Soc. Open Sci.* 2:150162. doi: 10.1098/rsos.150162

Bruggeman, H., Zosh, W., and Warren, W. (2007). Optic flow drives human visuo-locomotor adaptation. *Curr Biol. Cell* 17, 2035–2040. doi: 10.1016/j.cub.2007.10.059

Bruneau, J., Olivier, A.-H., and Pettré, J. (2015). Going through, going around: a study on individual avoidance of groups. *IEEE Trans. Vis. Comput. Graph.* 21, 520–528. doi: 10.1109/TVCG.2015.2391862

Burghardt, S., Seyfried, A., and Klingsch, W. (2013). Performance of stairs–fundamental diagram and topographical measurements. *Transp. Res. C* 37, 268–278. doi: 10.1016/j.trc.2013.05.002

Camazine, S. (2003). *Self-Organization in Biological Systems.* Princeton, NJ: Princeton University Press.

Carlson, J. M., Alderson, D. L., Stromberg, S. P., Bassett, D. S., Craparo, E. M., Guiterrez-Villarreal, F., et al. (2014). Measuring and modeling behavioral decision dynamics in collective evacuation. *PLoS ONE* 9:e87380. doi: 10.1371/journal.pone.0087380

Cassol, V. J., Testa, E. S., Jung, C. R., Usman, M., Faloutsos, P., Berseth, G., et al. (2017). Evaluating and optimizing evacuation plans for crowd egress. *IEEE Comput. Graph Appl.* 37, 60–71. doi: 10.1109/MCG.2017.3271454

Castellano, C., Fortunato, S., and Loreto, V. (2009). Statistical physics of social dynamics. *Rev. Mod. Phys. Am. Phys. Soc.* 81, 591–646. doi: 10.1103/RevModPhys.81.591

Chance, S. S., Gaunet, F., Beall, A. C., and Loomis, J. M. (1998). Locomotion mode affects the updating of objects encountered during travel: the contribution of vestibular and proprioceptive inputs to path integration. presence: teleoperators and virtual environments. *MIT Press* 7, 168–178.

Choi, J. W., Cho, S. H., Kim, Y. S., Kim, N. J., Kwon, S. S., and Shim, J. S. (2016). "A counting sensor for inbound and outbound people using IR-UWB radar sensors," in *2016 IEEE Sensors Applications Symposium (SAS),* 1–5.

Coşkun, A., Kara, A., Parlaktuna, M., Ozkan, M., and Parlaktuna, O. (2015). "People counting system by using kinect sensor," *2015 International Symposium on Innovations in Intelligent SysTems and Applications (INISTA),* 1–7.

Daamen, W., and Hoogendoorn, S. (2010). Capacity of doors during evacuation conditions. *Proc. Eng.* 3, 53–66. doi: 10.1016/j.proeng.2010.07.007

Dalal, N., and Triggs, B. (2005). "Histograms of oriented gradients for human detection," in *2005 CVPR 2005 IEEE Computer Society Conference on Computer, Vision, and Pattern, Recognition,* 886–893.

Damarla, T., Oispuu, M., Schikora, M., and Koch, W. (2016). "Tracking and counting multiple people using distributed seismic sensors," *2016 19th International Conference on Information Fusion (FUSION),* 1593–1599.

Davies, A. C., Yin, J. H., and Velastin, S. A. (1995). Crowd monitoring using image processing. *Electron. Commun. Eng. J.* 7, 37–47. doi: 10.1049/ecej: 19950106

Degond, P., Appert-Rolland, C., Moussaïd, M., Pettré J., and Theraulaz, G. (2013). A Hierarchy of heuristic-based models of crowd dynamics. *J. Stat. Phys.* 152, 1033–1068. doi: 10.1007/s10955-013-0805-x

Ding, H., Han, J., Liu, A. X., Zhao, J., Yang, P., Xi, W., et al. (2015). "Human object estimation via backscattered radio frequency signal," in *2015 IEEE Conference on Computer Communications (INFOCOM),* 1652–1660.

Drury, J., Cocking, C., Reicher, S., Burton, A., Schofield, D., Hardwick, A., et al. (2009). Cooperation versus competition in a mass emergency evacuation: a new laboratory simulation and a new theoretical model. *Behav. Res. Methods.* 41, 957–970. doi: 10.3758/BRM.41.3.957

Dutra, T. B., Marques, R., Cavalcante-Neto, J. B., Vidal, C. A., and Pettré, J. (2017). Gradient-based steering for vision-based crowd simulation algorithms. *Comput. Graph. Forum* 36, 337–348. doi: 10.1111/cgf.13130

Eagle, N., Pentland, A. S., and Lazer, D. (2009). Inferring friendship network structure by using mobile phone data. *Proc. Natl. Acad. Sci. U.S.A.* 106, 15274–15278. doi: 10.1073/pnas.0900282106

Epstein, R. A., Patai, E. Z., Julian, J. B., and Spiers, H. J. (2017). The cognitive map in humans: spatial navigation and beyond. *Nat. Neurosci.* 20, 1504–1513. doi: 10.1038/nn.4656

Epstein, R. A., and Vass, L. K. (2014). Neural systems for landmark-based wayfinding in humans. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369:20120533. doi: 10.1098/rstb.2012.0533

Fajen, B. R., and Warren, W. H. (2007). Behavioral dynamics of intercepting a moving target. *Exp. Brain Res.* 180, 303–319. doi: 10.1007/s00221-007-0859-6

Fink, P. W., Foo, P. S., and Warren, W. H. (2007). Obstacle avoidance during walking in real and virtual environments. *ACM Trans. Appl. Percept.* 4:2. doi: 10.1145/1227134.1227136

Foo, P., Warren, W. H., Duchon, A., and Tarr, M. J. (2005). Do humans integrate routes into a cognitive map? Map- versus landmark-based navigation of novel shortcuts. *J. Exp. Psychol. Learn. Mem. Cogn.* 31, 195–215. doi: 10.1037/0278-7393.31.2.195

Fukuzaki, Y., Mochizuki, M., Murao, K., and Nishio, N. A. (2014). "Pedestrian flow analysis system using wi-fi packet sensors to a real environment," in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication* (New York, NY: ACM), 721–730.

Gallistel, C. R. (1990). *The Organization of Learning.* Cambridge, MA: MIT press.

Gallup, A. C., Hale, J. J., Sumpter, D. J., Garnier, S., Kacelnik, A., Krebs, J., et al. (2012). Visual attention and the acquisition of information in human crowds. *Proc. Natl. Acad. Sci. U.S.A.* 109, 7245–7250. doi: 10.1073/pnas.1116141109

Ginsberg, J., Mohebbi, M., Patel, R., Brammer, L., Smolinski, M., and Brilliant, L. (2008). Detecting influenza epidemics using search engine query data. *Nature* 457, 1012–1014. doi: 10.1038/nature07634

Goel, S., Hofman, J. M., Lahaie, S., Pennock, D. M., and Watts, D. J. (2010). Predicting consumer behavior with Web search. *Proc. Natl. Acad. Sci. U.S.A.* 107, 17486–17490. doi: 10.1073/pnas.1005962107

González, M. C., Hidalgo, C. A., and Barabási, A.-L. (2008). Understanding individual human mobility patterns. *Nature* 453:779. doi: 10.1038/nature06958

Grübel, J., Thrash, T., Hölscher, C., and Schinazi, V. R. (2017). Evaluation of a conceptual framework for predicting navigation performance in virtual reality. *PLoS ONE* 12:e0184682. doi: 10.1371/journal.pone.0184682

Guy, S. J., Chhugani, J., Kim, C., Satish, N., Lin, M., Manocha, D., et al. (2009). "Clearpath: highly parallel collision avoidance for multi-agent simulation," in *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY: ACM), 177–187.

Haworth, B., Usman, M., Berseth, G., Khayatkhoei, M., Kapadia, M., and Faloutsos, P. (2017). CODE: Crowd-optimized design of environments. *Comput. Animat. Virtual Worlds* 28:e1749. doi: 10.1002/cav.1749

Helbing, D. (1992). "A mathematical model for behavioral changes by pair interactions," in *Economic Evolution and Demographic Change,* eds G. Haag, U. Mueller, and K.G. Troitzsch (Berlin; Heidelberg: Springer), 330–348.

Helbing, D., Brockmann, D., Chadefaux, T., Donnay, K., Blanke, U., Woolley-Meza, O., et al. (2014). Saving human lives: what complexity science and information systems can contribute. *J. Stat. Phys.* 158, 735–781. doi: 10.1007/s10955-014-1024-9

Helbing, D., Buzna, L., Johansson, A., and Werner, T. (2005). Self-organized pedestrian crowd dynamics: experiments, simulations, and design solutions. *Transport. Sci.* 39, 1–24. doi: 10.1287/trsc.1040.0108

Helbing, D., Johansson, A., and Al-Abideen, H. (2007). The dynamics of crowd disasters: an empirical study. *Phys. Rev. E* 75:46109. doi: 10.1103/PhysRevE.75.046109

Helbing, D., and Molnár, P. (1995). Social force model for pedestrian dynamics. *Phys. Rev. E* 51, 4282–4286.

Helbing, D., and Mukerji, P. (2012). Crowd disasters as systemic failures: analysis of the Love Parade disaster. *EPJ Data Sci.* 1, 1–40. doi: 10.1140/epjds7

Henderson, L. (1974). On the fluid mechanics of human crowd motion. *Transp. Res.* 8, 509–515.

Hodgson, E., Bachmann, E., and Waller, D. (2011). Redirected walking to explore virtual environments: assessing the potential for spatial interference. *ACM Trans. Appl. Percept.* 8:22. doi: 10.1145/2043603.2043604

Höllerer, T., Kuchera-Morin, J., and Amatriain, X. (2007). "The allosphere: a large-scale immersive surround-view instrument," in *Proceedings of the 2007 Workshop on Emerging Displays Technologies: Images and Beyond: The Future of Displays and Interacton* (New York, NY: ACM).

Hoogendoorn, S., and Daamen, W. (2005). Pedestrian behavior at bottlenecks. *Transp. Sci.* 39, 147–159. doi: 10.1287/trsc.1040.0102

Ioannou, C. C., Guttal, V., and Couzin, I. D. (2012). Predatory fish select for coordinated collective motion in virtual prey. *Science* 337, 1212–1215. doi: 10.1126/science.1218919

Jacobs, H. (1967). *To Count a Crowd. Columbia J Rev.* Columbia University, Graduate School of Journalism.

Jelić A., Appert-Rolland, C., Lemercier, S., and Pettré J. (2012). Properties of pedestrians walking in line: fundamental diagrams. *Phys. Rev. E Stat. Nonlin. Soft Matter. Phys.* 85:036111. doi: 10.1103/PhysRevE.85.036111

Kang, K., and Wang, X. (2014). Fully convolutional neural networks for crowd segmentation. *arXiv:1411.4464.*

Kannan, P. G., Venkatagiri, S. P., Chan, M. C., Ananda, A. L., and Peh, L.-S. (2012). "Low cost crowd counting using audio tones," in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems* (New York, NY: ACM), 155–168.

Khan, A. A., and Noji, E. K. (2016). Hajj stampede disaster, 2015: reflections from the frontlines. *Am. J. Disaster Med.* 11, 59–68. doi: 10.5055/ajdm.2016.0225

Khan, A. M., Siddiqi, M. H., and Lee, S.-W. (2013). Exploratory data analysis of acceleration signals to select light-weight and accurate features for real-time activity recognition on smartphones. *Sensors* 13, 13099–13122. doi: 10.3390/s131013099

Kinateder, M., and Warren, W. H. (2016). Social influence on evacuation behavior in real and virtual environments. *Front. Robot. AI* 3:43. doi: 10.3389/frobt.2016.00043

Kitchin, R., and Blades, M. (2002). *The Cognition of Geographic Space.* London: IB Tauris.

Klatzky, R. L., Loomis, J. M., Beall, A. C., Chance, S. S., and Golledge, R. G. (1998). Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychol. Sci.* 9, 293–298. doi: 10.1111/1467-9280.00058

Knapp, J. L. J. (2003). "Visual perception of egocentric distance in real and virtual environments," in *Virtual and Adaptive Environments* (Boca Raton, FL: CRC Press), 35–60. doi: 10.1201/9781410608888.pt1

Knierim, J. J., Kudrimoti, H. S., and McNaughton, B. L. (1998). Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. *J Neurophysiol.* 80, 425–446. doi: 10.1152/jn.1998.80.1.425

Kretz, T., Grünebohm, A., and Schreckenberg, M. (2006). Experimental study of pedestrian flow through a bottleneck. *J. Stat. Mech.* 2006:P10014. doi: 10.1088/1742-5468/2006/10/P10014

Lapointe, J.-F., Savard, P., and Vinson, N. G. (2011). A comparative study of four input devices for desktop virtual walkthroughs. *Comput. Hum. Behav.* 27, 2186–2191. doi: 10.1016/j.chb.2011.06.014

Liang, R., Zhu, Y., and Wang, H. (2014). Counting crowd flow based on feature points. *Neurocomputing* 133, 377–384. doi: 10.1016/j.neucom.2013.12.040

Liddle, J., Seyfried, A., Klingsch, W., Rupprecht, T., Schadschneider, A., Winkens, A., et al. (2009). An experimental study of pedestrian congestions: influence of bottleneck width and Length. *Traffic Granular.* arXiv:0911.4350.

Lin, S.-F., Chen, J.-Y., and Chao, H.-X. (2001). "Estimation of number of people in crowded scenes using perspective transformation," in *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 645–654.

Liu, L., Kuo, S. M., and Zhou, M. (2009). "Virtual sensing techniques and their applications,"in 2009 *International Conference on Networking, Sensing and Control,* 31–36.

Louison, C., Ferlay, F., and Mestre, D. R. (2017). "Spatialized vibrotactile feedback contributes to goal-directed movements in cluttered virtual environments," in *2017 IEEE Symposium on 3D User Interfaces (3DUI),* 99–102.

Ma, R., Li, L., Huang, W., Tian, Q. (2004). "On pixel count based crowd density estimation for visual surveillance," in *IEEE Conference on Cybernetics and Intelligent Systems,* 170–173.

Mao, A., Mason, W., Suri, S., and Watts, D. J. (2016). An experimental study of team size and performance on a complex task. *PLoS ONE* 11:e0153048. doi: 10.1371/journal.pone.0153048

Marana, A. N., Cavenaghi, M. A., Ulson, R. S., and Drumond, F. L. (2005). *Real-Time Crowd Density Estimation Using Images. Advances in Visual Computing.* Berlin; Heidelberg: Springer.

Mason, W., and Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. *Behav. Res. methods* 44, 1–23. doi: 10.3758/s13428-011-0124-6

Mason, W., and Watts, D. (2012). Collaborative learning in networks. *Proc. Natl. Acad. Sci. U.S.A.* 109, 764–769. doi: 10.1073/pnas.1110069108

Molka-Danielsen, J., and Chabada, M. (2010). "Application of the 3D multi user virtual environment of second life to emergency evacuation simulation," in *2010 43rd Hawaii International Conference on System Sciences*, 1–9.

Montello, D. R. (2005). *Navigation.* Cambridge: Cambridge University Press.

Mori, T., Suemasu, Y., Noguchi, H., and Sato, T. (2004). "Multiple people tracking by integrating distributed floor pressure sensors and RFID system," in *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat No04CH37583),* 5271–5278.

Moussaïd, M., Brighton, H., and Gaissmaier, W. (2015). The amplification of risk in experimental diffusion chains. *Proc. Natl. Acad. Sci. U.S.A.* 112, 5631–5636. doi: 10.1073/pnas.1421883112

Moussaïd, M., Garnier, S., Theraulaz, G., and Helbing, D. (2009). *Collective Information Processing and Pattern Formation in Swarms, Flocks, and Crowds.* Top. Cogn. Sci. 1, 1–29. doi: 10.1111/j.1756-8765.2009.01028.x

Moussaïd, M., Guillot, E. G., Moreau, M., Fehrenbach, J., Chabiron, O., Lemercier, S., et al. (2012). Traffic instabilities in self-organized pedestrian crowds. *PLoS Comput. Biol.* 8:e1002442. doi: 10.1371/journal.pcbi.1002442

Moussaïd, M., Helbing, D., Garnier, S., Johansson, A., Combe, M., and Theraulaz, G. (2009). Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proc. R. Soc. B Biol. Sci.* 276, 2755–2762. doi: 10.1098/rspb.2009.0405

Moussaïd, M., Helbing, D., and Theraulaz, G. (2011). How simple rules determine pedestrian behavior and crowd disasters. *Proc. Natl. Acad. Sci. U.S.A.* 108, 6884–6888. doi: 10.1073/pnas.1016507108

Moussaïd, M., Herzog, S. M., Kämmer, J. E., and Hertwig, R. (2017). Reach and speed of judgment propagation in the laboratory. *Proc. Natl. Acad. Sci. U.S.A.* 114, 4117–4122. doi: 10.1073/pnas.1611998114

Moussaïd, M., Kapadia, M., Thrash, T., Sumner, R. W., Gross, M., Helbing, D., et al. (2016). Crowd behaviour during high-stress evacuations in an immersive virtual environment. *J. R. Soc. Interface R. Soc.* 13:20160414. doi: 10.1098/rsif.2016.0414

Musa, A. B. M., and Eriksson, J. (2012). "Tracking unmodified smartphones using wi-fi monitors," in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems* (New York, NY: ACM), 281–294.

Nagy, M., Ákos, Z., Biro, D., and Vicsek, T. (2010). Hierarchical group dynamics in pigeon flocks. *Nature* 464, 890–893. doi: 10.1038/nature08891

Narang, S., Best, A., Feng, A., and Kang, S. (2017a). *Motion Recognition of Self and Others on Realistic 3D Avatars and Virtual Worlds.* Wiley Online Library. Available online at: http://onlinelibrary.wiley.com/doi/10.1002/cav.1762/full

Narang, S., Best, A., Shapiro, A., and Manocha, D. (2017b). "Generating virtual avatars with personalized walking gaits using commodity hardware," in *Proceedings of the on Thematic Workshops of ACM Multimedia* (New York, NY: ACM), 219–227.

Nescher, T., Huang, Y.-Y., and Kunz, A. (2014). "Planning redirection techniques for optimal free walking experience using model predictive control," in *2014 IEEE Symposium on 3D User Interfaces (3DUI)*, 111–118.

Normoyle, A., Drake, J., and Safonova, A. (2012). *Egress Online: Towards leveraging Massively, Multiplayer Environments for Evacuation Studies.* Citeseer. Available online at: https://repository.upenn.edu/cgi/viewcontent.cgi?article=2023&context=cis_reports

Olivier, A.-H., Marin, A., Crétual, A., Berthoz, A., and Pettré, J. (2013). Collision avoidance between two walkers: role-dependent strategies. *Gait Posture* 38, 751–756. doi: 10.1016/j.gaitpost.2013.03.017

Ondrej, J., Pettré J., Olivier, A.-H., and Donikian, S. (2010). A synthetic-vision based steering approach for crowd simulation. *ACM Trans. Graph.* 29, 1–9. doi: 10.1145/1778765.177886

Ouyang, W., and Wang, X. (2012). "A discriminative deep model for pedestrian detection with occlusion handling," in *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3258–3265.

Ouyang, W., and Wang, X. (2013). "Joint deep learning for pedestrian detection," in *2013 IEEE International Conference on Computer Vision (ICCV),* 2056–2063.

Pérez-Escudero, A., Vicente-Page, J., Hinz, R. C., Arganda, S., and De Polavieja, G. G. (2014). idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nat. Methods* 11, 743–748. doi: 10.1038/nmeth.2994

Pražák, M., and O'Sullivan, C. (2011). "Perceiving human motion variety," in Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization (New York, NY: ACM), 87–92.

Presson, C. C., and Montello, D. R. (1988). Points of reference in spatial cognition: stalking the elusive landmark. *Br. J. Dev. Psychol.* 6, 378–381.

Qiao, G., Yoon, S., Kapadia, M., and Pavlovic, V. (2017). The role of data-driven priors in multi-agent crowd trajectory estimation. *arXiv.*

Raybould, M., Mules, T., Fredline, E., and Tomljenovic, R. (2000). Counting the herd. using aerial photography to estimate attendance at open events. *Event Manage.* 6, 25–32.

Razzaque, S, Kohn, Z., Whitton, M. C. (2001). "Redirected walking," in *Proceedings of Eurographics,* 105–106.

Richardson, A. E., Montello, D. R., and Hegarty, M. (1999). Spatial knowledge acquisition from maps and from navigation in real and virtual environments. *Mem. Cogn.* 27, 741–750. doi: 10.3758/BF03211566

Richardson, A. R., and Waller, D. (2005). The effect of feedback training on distance estimation in virtual environments. *Appl. Cogn. Psychol.* 19, 1089–1108. doi: 10.1002/acp.1140

Rittscher, J., Kelliher, T., and Tu, P. (2005). *System and Method for Segmenting Crowded Environments into Individual Objects.* U.S. Patent Application No. 10/942,056.

Ryu, J., and Kim, G. J. (2004). "Using a vibro-tactile display for enhanced collision perception and presence," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology.* (New York, NY: ACM), 89–96.

Schadschneider, A., Klingsch, W., Klüpfel, H., Kretz, T., Rogsch, C., and Seyfried, A. (2011). "Evacuation dynamics: empirical results, modeling and applications,"

in *Extreme Environmental Events,* ed R. Meyers (New York, NY: Springer), 517–550.

Schauer, L., Werner, M., and Marcus, P. (2014). "Estimating crowd densities and pedestrian flows using wi-fi and bluetooth," in *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services* (Belgium: ICST Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 171–177.

Serafin, S., Nordahl, R., De Götzen, A., Erkut, C., Geronazzo, M., and Avanzini, F. (2015). "Sonic interaction in virtual environments," in *2015 IEEE 2nd VR Workshop on Sonic Interactions for Virtual Environments (SIVE),* 1–2.

Sermanet, P., Kavukcuoglu, K., Chintala, S., and LeCun, Y. (2013). "Pedestrian detection with unsupervised multi-stage feature learning," in *2013 IEEE Conference, Computer Vision and Pattern Recognition (CVPR),* 3626–3633.

Seyfried, A., Passon, O., Steffen, B., Boltes, M., Rupprecht, T., and Klingsch, W. (2009). New insights into pedestrian flow through bottlenecks. *Transport. Sci.* 43, 395–406. doi: 10.1287/trsc.1090.0263

Shao, J., Kang, K., Change Loy, C., Wang, X. (2015). "Deeply learned attributes for crowded scene understanding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4657–4666.

Shao, J., Loy, C. C., Kang, K., and Wang, X. (2017). "Crowded scene understanding by deeply learned volumetric slices," in *IEEE Trans Circuits Syst Video Technol,* 27, 613–623.

Sindagi, V. A., and Patel, V. M. (2017). "Generating high-quality crowd density maps using contextual pyramid cnns," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 1879–1888.

Singh, S., Kapadia, M., Hewlett, B., Reinman, G., and Faloutsos, P. (2011a). "A modular framework for adaptive agent-based steering," in *Symposium on Interactive 3D Graphics and Games* (New York, NY: ACM), 141–150.

Singh, S., Kapadia, M., Reinman, G., and Faloutsos, P. (2011b). "Footstep navigation for dynamic crowds," in *Symposium on Interactive 3D Graphics and Games* (New York, NY: ACM), 203.

Souman, J. L., Giordano, P. R., Frissen, I., Luca, A. D., and Ernst, M. O. (2010). Making virtual walking real: perceptual evaluation of a new treadmill control algorithm. *ACM Trans. Appl. Percept.* 7:11. doi: 10.1145/1670671.16 70675

Still, G. K. (2000). *Crowd Dynamics*, Doctoral dissertation, University of Warwick.

Stowers, J. R., Hofbauer, M., Bastien, R., Griessner, J., Higgins, P., Farooqui, S., et al. (2017). Virtual reality for freely moving animals. *Nat. Methods* 14, 995–1002. doi: 10.1038/nmeth.4399

Strandburg-Peshkin, A., Farine, D. R., Couzin, I. D., and Crofoot, M. C. (2015). Shared decision-making drives collective movement in wild baboons. *Science* 348, 1358–1361. doi: 10.1126/science.aaa5099

Strandburg-Peshkin, A., Farine, D. R., Crofoot, M. C., and Couzin, I. D. (2017). Habitat and social factors shape individual decisions and emergent group structure during baboon collective movement. *eLife Sci.* 6:e19505. doi: 10.7554/eLife.19505

Taube, J. S., Valerio, S., and Yoder, R. M. (2013). Is navigation in virtual reality with FMRI really navigation? *J. Cogn. Neurosci.* 25, 1008–1019. doi: 10.1162/jocn_a_00386

Thrash, T., Kapadia, M., Moussaid, M., Wilhelm, C., Helbing, D., Sumner, R. W., et al. (2015). Evaluation of control interfaces for desktop virtual environments. presence: teleoperators and virtual environments. *MIT Press* 24, 322–334. doi: 10.1162/PRES_a_00237

Tian, Y., Luo, P., Wang, X., Tang, X. (2015). "Deep learning strong parts for pedestrian detection," in *Proceedings of the IEEE International Conference on Computer Vision*, 1904–1912.

van den Berg, J., Guy, S. J., Lin, M., Manocha, D. (2011). "Reciprocal n-body collision avoidance," in *Robotics Research*, eds C. Pradalier, R. Siegwart, and G. Hirzinger (Berlin; Heidelberg: Springer-Verlag), 3–19.

Van Diggelen (2009). *Van Diggelen FST. A-GPS: Assisted GPS, GNSS, and SBAS.* Norwood: Artech House.

Vassallo, C., Olivier, A.-H., Souères, P., Crétual, A., Stasse, O., and Pettré, J. (2017). How do walkers avoid a mobile robot crossing their way? *Gait Posture.* 51, 97–103. doi: 10.1016/j.gaitpost.2016.09.022

Wagoum, A. U., Tordeux, A., and Liao, W. (2017). Understanding human queuing behaviour at exits: an empirical study. *R. Soc. Open Sci.* 4:160896. doi: 10.1098/rsos.160896

Waller, D., and Lippa, Y. (2007). Landmarks as beacons and associative cues: their role in route learning. *Mem. Cogn.* 35, 910–924. doi: 10.3758/BF03193465

Waller, D., Loomis, J. M., and Haun, D. B. M. (2004). Body-based senses enhance knowledge of directions in large-scale environments. *Psychon. Bull. Rev.* 11, 157–163. doi: 10.3758/BF03206476

Warren, W. H., and Fajen, B. R. (2004). Behavioral dynamics of human locomotion. *Ecol Psychol.* 16, 61–66. doi: 10.1207/s15326969eco1601_8

Warren, W. H., and Fajen, B. R. (2008). "Behavioral dynamics of visually guided locomotion," in *Coordination: Neural, Behavioral and Social Dynamics*, eds A. Fuchs and V. K. Jirsa (Berlin; Heidelberg: Springer Berlin Heidelberg), 45–75.

Weppner, J., Bischke, B., and Lukowicz, P. (2016). "Monitoring crowd condition in public spaces by tracking mobile consumer devices with wifi interface," in *Proceedings of the* 2016 *ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct* (New York, NY: ACM), 1363–1371.

Wirz, M., Franke, T., Roggen, D., Mitleton-Kelly, E., Lukowicz, P., Tröster, G., et al. (2012). "Infrastructure for collaborative enterprises," in *2012 IEEE 21st International Workshop on Enabling Technologies*, 367–372.

Wirz, M., Franke, T., Roggen, D., Mitleton-Kelly, E., Lukowicz, P., and Tröster, G. (2013). Probing crowd density through smartphones in city-scale mass gatherings. *EPJ Data Sci.* 2:5. doi: 10.1140/epjds17

Nonapeptide Receptor Distributions in Promising Avian Models for the Neuroecology of Flocking

Naomi R. Ondrasek[1]*, Sara M. Freeman[2,3], Karen L. Bales[2,3] and Rebecca M. Calisi[1]

[1] Department of Neurobiology, Physiology and Behavior, University of California, Davis, Davis, CA, United States,
[2] Department of Psychology, University of California, Davis, Davis, CA, United States, [3] California National Primate Research Center, University of California, Davis, Davis, CA, United States

Collective behaviors, including flocking and group vocalizing, are readily observable across a diversity of free-living avian populations, yet we know little about how neural and ecological factors interactively regulate these behaviors. Because of their involvement in mediating a variety of social behaviors, including avian flocking, nonapeptides are likely mediators of collective behaviors. To advance the neuroecological study of collective behaviors in birds, we sought to map the neuroanatomical distributions of nonapeptide receptors in three promising avian models that are found across a diversity of environments and widely ranging ecological conditions: European starlings, house sparrows, and rock doves. We performed receptor autoradiography using the commercially available nonapeptide receptor radioligands, ^{125}I-ornithine vasotocin analog and ^{125}I-linear vasopressin antagonist, on brain tissue sections from wild-caught individuals from each species. Because there is known pharmacological cross-reactivity between nonapeptide receptor subtypes, we also performed a novel, competitive-binding experiment to examine the composition of receptor populations. We detected binding in numerous regions throughout the brains of each species, with several similarities and differences worth noting. Specifically, we report that all three species exhibit binding in the lateral septum, a key brain area known to regulate avian flocking. In addition, sparrows and starlings show dense binding in the dorsal arcopallium, an area that has received scant attention in the study of social grouping. Furthermore, our competitive binding results suggest that receptor populations in sparrows and starlings differ in the lateral septum versus the dorsal arcopallium. By providing the first comprehensive maps of nonapeptide receptors in European starlings, house sparrows, and rock doves, our work supports the future use of these species as avian models for neuroecological studies of collective behaviors in wild birds.

Keywords: neuroecology, oxytocin, vasopressin, mesotocin, vasotocin, grouping behavior

INTRODUCTION

Diverse examples of collective behaviors exist across the animal kingdom, but perhaps most conspicuous is the formation of large, coordinated groups in which individuals communicate, move, and forage together (Parrish and Edelstein-Keshet, 1999). The ecological pressures that drive or stabilize the evolution of these groups have been considered in depth (e.g., Alexander, 1974;

Emlen, 1982; Solomon, 2003), but we know very little about the neural processes that prompt individuals to participate in these aggregations. Free-living birds are ideal for investigating the emergence of collective behaviors from interactions among neural systems and ecological factors—the focus of an emerging field called neuroecology (Sherry, 2006; Zimmer and Derby, 2011)—because they frequently form conspicuous groups that are comprised of individuals that feed, evade predators, and vocalize together (Helm et al., 2006). However, the neuroecology of collective behaviors has received little attention, perhaps in part because we lack well-developed organismal models suited to these types of investigations. We sought to address this gap by taking the first steps toward developing three globally distributed avian species—house sparrows (*Passer domesticus*), European starlings (*Sturnus vulgaris*), and rock doves (*Columba livia*)—as potential models for neuroecological studies of collective behaviors.

Because of their ability to invade, inhabit, and form groups in a diversity of environments, house sparrows, European starlings, and rock doves are particularly advantageous for studying how ecological variations influence the neural processes underlying collective behaviors. Since their introductions via the eastern coast of North America, these species have spread across vast swaths of the continent and today, members of each species number in the millions throughout the United States. Because of their wide distributions, these species are found across a spectrum of environmental conditions, including a variety of climates, urbanization gradients, and ecological communities (Cabe, 1993; Johnston and Janiga, 1995; Clergeau et al., 1998; Anderson, 2006). Thus, sparrows, starlings, and rock doves are ideal for intraspecies, inter-population comparisons that can reveal much about the impacts of varying ecological factors on the neurobiology underlying collective behaviors.

In addition to selecting ideal avian models, advancing the neuroecological study of collective behaviors requires that we identify candidate neural systems, ideally with demonstrated involvement in regulating social behaviors. The nonapeptide (NP) systems are an excellent place to start because they mediate a wide variety of social behaviors, including pair bonding, parent-offspring bonding, same-sex interactions, and group size preference (reviewed in Beery et al., 2016). All vertebrate species examined thus far produce NPs, a highly conserved class of neurohormones that includes oxytocin, vasopressin, and their non-mammalian homologs mesotocin and vasotocin, respectively (Gimpl and Fahrenholz, 2001; Goodson, 2005, 2013). Thus, discoveries made regarding the role of NP systems in avian collective behaviors can provide insights that support and guide similar investigations in other animal groups.

One limitation for examining NP system function in house sparrows, European starlings, and rock doves is that NP receptors have never been mapped in these species. Such maps are necessary complements to laboratory investigations, which in turn are needed to demonstrate causal links between neural and behavioral processes. In addition, studying NP systems is challenging due to a high level of structural homology and pharmacological cross-reactivity among the four subtypes of NP receptors (Acher et al., 1995; Ocampo Daza et al., 2012). This characteristic has made it difficult to identify the specific

functional contributions of each receptor subtype to behavior, particularly in birds (Leung et al., 2009). Phylogenetic analyses of receptor amino acid sequences in a handful of avian species have identified four avian NP receptor subtypes (summarized in Leung et al., 2011). These studies have also shown that the two subtypes that are most highly expressed in the avian brain are vasotocin (VT) receptor 4 (VT4), which has a high degree of sequence homology to the mammalian vasopressin receptor 1a (V1aR) (Leung et al., 2011; Genbank ACCN abv24997), and avian VT3, which shares a high sequence identity with the mammalian oxytocin receptor (OTR) (Gubrij et al., 2005). Thus, our investigation focused on identifying VT4 (referred to here as V1aR-like) and VT3 (referred to here as OTR-like) as the relevant NP receptors for the current study.

To address these challenges and further the development of promising avian models for the study of the neuroecology of flocking behavior, we sought to accomplish two goals: first, to map the distribution of NP receptors in brain tissue from house sparrows, European starlings, and rock doves, and second, to identify potentially heterogenous populations of NP receptors in these species. To this end, we performed receptor autoradiography using two radioligands that are commonly employed in studies of mammalian NP receptors: ^{125}I-ornithine vasotocin analog (^{125}I-OVTA), which is used to label OTR, and ^{125}I-linearized vasopressin antagonist (^{125}I-LVA), which is used to label V1aR. We expected that ^{125}I-OVTA and ^{125}I-LVA would primarily label VT3 (OTR-like) and VT4 (V1aR-like), respectively. However, these radioligands produce overlapping patterns of binding in the brains of other avian species (Leung et al., 2009), which may suggest that these molecular tools bind more promiscuously to the avian NP receptors than they do in rodents. Alternatively, such overlap in radioligand binding may reflect true mixed receptor populations in specific regions of the avian brain.

To examine which specific receptor subtypes contribute to the binding patterns of each radioligand, we performed a competitive binding experiment to assess the impact that a V1aR competitor, the Manning compound, would have on ^{125}I-OVTA and ^{125}I-LVA binding. Due to its strong affinity for V1aR, the Manning compound is frequently used in studies of mammalian NP systems, both as a competitor to distinguish among different receptor classes for mapping purposes, and as an antagonist to examine V1aR contributions to behavioral regulation (Manning et al., 2012). We placed particular focus on determining how the Manning compound impacts ^{125}I-OVTA and ^{125}I-LVA binding in the lateral septum (LS) because NP receptors have been identified in this region in several avian species, and the LS has been implicated in the regulation of avian flocking behaviors (Goodson et al., 2009b; Leung et al., 2009; Kelly et al., 2011).

We selected the Manning compound for use as a putative competitor for the avian V1aR-like receptor (VT4) after first considering the molecular basis for our hypothesized pharmacological homology. In mammalian systems, the amino acids in the third and eighth positions for endogenous NPs are known to confer ligand-binding specificity by interacting with specific amino acid residues in V1aR and OTR; specifically,

amino acid residues 509 and 609 in V1aR interact with the third amino acid in vasopressin (Chini et al., 1996), and residue 115 in V1aR interacts with the eighth amino acid in vasopressin (Chini et al., 1995). These three key amino acid residues in V1aR, which confer binding specificity to vasopressin, are identical in the amino acid sequence of avian VT4 (Leung et al., 2011). Additionally, the Manning compound and vasopressin are also identical in the amino acids present at the third and eighth positions (Kruszynski et al., 1980); thus, we expected that the Manning compound should bind selectively to VT4, the putative V1aR-like avian NP receptor.

Multiple studies across several avian species demonstrate that ^{125}I-OVTA binds to multiple brain areas, whereas ^{125}I-LVA only produces visible labeling in some, but not all, species (Goodson et al., 2006; Leung et al., 2009). We predicted that we would observe similar trends in this experiment; specifically, that ^{125}I-OVTA would label NP receptors in all three of our examined species, while ^{125}I-LVA would bind to receptors in only a subset of these species, across fewer brain regions, or at lower levels compared to ^{125}I-OVTA. We further predicted that the Manning compound would produce more radioligand displacement in the LS when labeled receptors are V1aR-like; specifically, we expected that the Manning compound would displace ^{125}I-LVA more than ^{125}I-OVTA, if these radioligands are binding selectively to their corresponding avian NP receptors. Alternatively, V1aR-like and OTR-like receptors in these species may bind ^{125}I-LVA, ^{125}I-OVTA, and the Manning compound with similar affinities; if this is the case, we expected that the Manning compound would displace ^{125}I-LVA and ^{125}I-OVTA to a similar degree.

MATERIALS AND METHODS

Animals

All birds were free-living and captured using mist nets or clap traps between 2013 and 2016. Specifically, male house sparrows ($n = 3$) were captured in November 2014 in Davis, CA, United States; female European starlings ($n = 3$) were captured in Tracy, CA, United States in January 2014; and female rock doves ($n = 3$) were captured either in Tracy, CA, United States (2 individuals) in September 2013, or in Davis, CA, United States (1 individual) in April 2016. Animal procedures were approved by the Animal Care and Use Committee of the University of California, Davis and abided by federal and state guidelines for animal care and use.

Tissue Collection and Preparation

After capture, birds were rapidly anesthetized under isoflurane and decapitated. Brains were removed, frozen immediately on dry ice, and transferred to $-80°C$ for storage until coronal sectioning on a cryostat. Brains were sectioned at 20 μm increments into 4 adjacent series at $-20°C$ and subsequently mounted on to Fisher Superfrost plus slides (Fisher, Pittsburg, PA, United States), which were stored in sealed slide boxes and returned to $-80°C$ until use for receptor autoradiography.

Receptor Autoradiography for NP Receptors

Nonapeptide receptor autoradiography assays were carried out as previously described (Perkeybile et al., 2015; Guoynes et al., 2018; Hartman et al., 2018). Sections were allowed to thaw in slide boxes for 1 h at room temperature and then placed in racks to dry. Slides were fully submerged in 0.1% paraformaldehyde, followed by two washes in 50 mM Tris-HCl (pH 7.4). Slides were then incubated for 1 h in a solution of 50 mM Tris-HCl (pH 7.4) with 10 mM MgCl$_2$, 0.1% bovine serum albumin, and 50 pM of radioligand. In this binding step, each series was then incubated in one of the following radioligand conditions: 50 pM of the OTR radioligand, ^{125}I-OVTA (PerkinElmer, Boston, MA, United States) or 50 pM of the V1aR radioligand, ^{125}I-LVA (PerkinElmer, Boston, MA, United States). Two of these series were incubated either in 50 pM ^{125}I-OVTA plus 1 μM of the highly selective V1aR antagonist, the Manning compound, or 50 pM ^{125}I-LVA plus 1 μM of the Manning compound. After the incubation period, slides were washed in multiple changes of chilled 50 mM Tris-HCl (pH 7.4) with 10 mM MgCl$_2$. Slides were then placed in a final rinse of this solution for 30 min, with gentle stirring, then rinsed in ddH$_2$O and allowed to air-dry overnight. Slides were then apposed to Carestream BioMax MR film (Kodak, Rochester, NY, United States) with a set of ten ^{125}I microscale standards (American Radiolabeled Chemicals, Inc., St. Louis, MO, United States) for 4 days, then developed and analyzed.

Imaging and Quantification

Photography of autoradiography films and quantification of regions with visible binding above background were accomplished using the MCID Digital Densitometry Core System (Interfocus Imaging, Cambridge, United Kingdom). Optical binding density (OBD) was quantified by extrapolation from a standard curve, which was constructed using a set of autoradiography standards (American Radiolabeled Chemicals, Inc., St. Louis, MO, United States) that were apposed to film in conjunction with specimen slides. For each bird, specific binding values for areas with visible binding were averaged across three sections for each area of interest. To account for individual differences in non-specific binding, OBD was measured in each section in a background area where no visible binding was apparent. For each section, specific binding was calculated by subtracting the non-specific binding value from OBD values obtained for each area. For the competitive binding experiment, labeling in the LS was quantified across all three species, and labeling in the dorsal arcopallium (Ad) was quantified in house sparrows and starlings, but not in rock doves due to a lack of ^{125}I-OVTA and ^{125}I-LVA binding in this area.

Identification of labeled brain regions was accomplished by referencing avian brain atlases and key neuroanatomical landmarks, visible on slides and in photomicrographs. Brain regions were identified in house sparrows using Nixdorf-Bergweiler and Bischof (2007), in European starlings using Nixdorf-Bergweiler and Bischof (2007) and De Groof et al. (2016), and in rock doves using Karten and Hodos (1967). Names

for brain regions identified using Karten and Hodos (1967) were updated according to Reiner et al. (2004) and Jarvis et al. (2005).

Statistical Analysis

Because work in several songbird species implicates the LS in the regulation of flocking behavior, comparisons of OBD values across binding conditions were planned *a priori* and used two-tailed *t*-tests or Wilcoxon rank sum tests. To minimize the risk of inflating the type 1 error rate, only a subset of all possible comparisons was performed. These planned comparisons excluded only those that present little heuristic value, an approach that has been described in detail elsewhere (Ruxton and Beauchamp, 2008; Ondrasek et al., 2015). Specifically, the comparisons were as follows: ^{125}I-OVTA versus ^{125}I-LVA, to identify differences in binding density for these two ligands; ^{125}I-OVTA versus ^{125}I-OVTA + Manning compound, to examine the impacts of the competitor on ^{125}I-OVTA binding; ^{125}I-LVA versus ^{125}I-LVA + Manning compound, to assess the impacts of the competitor on ^{125}I-LVA binding; and ^{125}I-LVA + Manning compound versus ^{125}I-OVTA + Manning compound, to determine if the competitor had differential impacts on ^{125}I-LVA versus ^{125}I-OVTA binding. The decision to use t-tests assuming equal or unequal variances was made subsequent to Bartlett's test for unequal variance.

In starlings and house sparrows, ^{125}I-LVA and ^{125}I-OVTA binding densities were particularly high in Ad. Because dense ^{125}I-OVTA labeling in the arcopallium has been reported in several other songbird species (Leung et al., 2009; Wilson et al., 2016), and because of this area's putative homology to the mammalian amygdala—a region with significant contributions to social behavior in mammals (Jarvis, 2009; Hanics et al., 2016)—*post hoc* comparisons of OBD values across binding conditions were performed using the Steel-Dwass method for non-parametric multiple comparisons, following ANOVA.

To provide a further test of differences in ^{125}I-OVTA and ^{125}I-LVA binding, and to identify general species effects on radioligand binding patterns, we combined ^{125}I-OVTA and ^{125}I-LVA optical binding densities for all three species and performed a principal component (PC) analysis. Only the 32 brain areas showing either ^{125}I-OVTA or ^{125}I-LVA binding in at least one species were included in the analysis. PC scores were subsequently analyzed using ANOVAs and non-parametric tests. Additional details regarding these analyses—including PC loadings, statistical test outcomes, and an interpretation of the results— may be found in the **Supplementary Materials**.

All statistical analyses were completed using JMP Pro 12 (SAS Institute Inc., Cary, NC, United States). Means ± SEM are reported throughout, and all OBD values included in statistical analyses and reported in figures and tables have been corrected for non-specific binding as described above.

RESULTS

General Observations

We expected that ^{125}I-OVTA and ^{125}I-LVA would primarily label the avian OTR-like receptor (VT3) and the avian V1aR-like receptor (VT4), respectively. Based on analyses of homologous amino acid sequences, we also hypothesized that the Manning compound should bind selectively to V1aR-like receptors in avian brains, as it does in mammals. Binding of ^{125}I-OVTA and ^{125}I-LVA was widely dispersed across a variety of brain regions in European starlings and rock doves, but comparatively more restricted in house sparrows (for complete lists and abbreviations of regions showing binding, see **Tables 1**–**4**). Save one exception in starling brain, ^{125}I-LVA binding always occurred in areas that also showed ^{125}I-OVTA binding, but the reverse was not always true. For instance, in rock doves, ^{125}I-OVTA, but not ^{125}I-LVA, signal was apparent in arcopallium (A), basorostral pallial nucleus (Bas), entopallium (E), and mesopallium (M). Of the two radioligands, ^{125}I-OVTA binding produced a more

TABLE 1 | Abbreviations for avian brain areas.

Abbreviation	Brain region
A	Arcopallium
Ad	Dorsal arcopallium
Al	Lateral arcopallium
APH	Parahippocampal area
Bas	Basorostral pallial nucleus
Cb	Cerebellum
CcS	Caudocentral septum
CMM	Caudomedial mesopallium
CoS	Commissural septal nucleus
H	Hyperpallium
LMAN	Lateral magnocellular nucleus of the anterior nidopallium
LS	Lateral septum
M	Mesopallium
MBH	Mediobasal hypothalamus
MMAN	Medial magnocellular nucleus of the anterior nidopallium
MSt	Medial striatum
N	Nidopallium
NCM	Caudal medial nidopallium
OMd	Dorsal nucleus of the oculomotor nerve
OMv	Ventral nucleus of the oculomotor nerve
Ov	Ovoid nucleus
pHVC	Para-high vocal center
RA	Robust nucleus of the arcopallium
TeO	Optic tectum
TnA	Nucleus taeniae of the amygdala
Uva	Uvaeform nucleus
VMH	Ventromedial hypothalamus
DLP	Dorsolateral nucleus of the posterior thalamus
DMP	Dorsomedial nucleus of the posterior thalamus
E	Entopallium
Hp	Hippocampus
Lhy	Lateral hypothalamus
MVL	Ventrolateral nucleus of the mesopallium
NIM	Intermediate medial nidopallium
SGP	Periventricular gray and fibrous tectal layers

intense signal in most areas, whereas incubation with ^{125}I-LVA resulted in far more non-specific binding (i.e., unilateral binding, often without distinct shape or edges, that was not repeated across two or more sections). In addition, intraspecies variation in ^{125}I-OVTA and ^{125}I-LVA binding was apparent in numerous brain areas across all three species; for many regions, binding was observed in some, but not all individuals (**Tables 2–4**).

TABLE 2 | Mean optical binding density (±SEM) of ^{125}I-OVTA and ^{125}I-LVA in brain regions with visible binding in house sparrow.

Brain region	^{125}I-OVTA	^{125}I-LVA
Ad	3433 ± 658(*ABC*)	2956 ± 195(*AB*)
APH	1153 ± 221(*ABC*)	538 ± 92(*BC*)
CMM	482 ± 28(*B*)	245 ± 23(*B*)
H	429 ± 71(*B*)	360 ± 41(*C*)
LS	1087 ± 67(*ABC*)	158 ± 157(*C*)
MMAN	732 ± 158(*ABC*)	369 ± 120(*B*)
MSt	240 ± 39(*C*)	693 ± 84(*AC*)
OMd	911 ± 114(*B*)	651 ± 54(*B*)
OMv	793 ± 19(*B*)	797 ± 107(*B*)
pHVC	1290 ± 232(*AB*)	815 ± 102(*B*)
TeO	571 ± 55(*ABC*)	636 ± 81(*ABC*)

Letters (A, B, or C) represent individual birds that showed binding in the indicated brain area.

TABLE 3 | Mean optical binding density (±SEM) of ^{125}I-OVTA and ^{125}I-LVA in brain regions with visible binding in European starling.

Brain region	^{125}I-OVTA	^{125}I-LVA
Ad	2685 ± 85(*ABC*)	1481 ± 70(*ABC*)
Al	3976 ± 186(*ABC*)	2587 ± 269(*ABC*)
APH	968 ± 117(*ABC*)	627 ± 93(*ABC*)
Bas	1595 ± 45(*C*)	–
CcS	1070 ± 168(*ABC*)	633 ± 111(*BC*)
CMM	–	630 ± 32(*C*)
CoS	2190 ± 150(*AC*)	–
H	768 ± 127(*ABC*)	765 ± 215(*C*)
LMAN	981 ± 139(*ABC*)	–
LS	4234 ± 490(*ABC*)	1480 ± 456(*ABC*)
M	983 ± 122(*ABC*)	841 ± 107(*BC*)
MBH	595 ± 233(*B*)	–
MSt	591 ± 91(*B*)	–
N	1026 ± 140(*ABC*)	855 ± 182(*BC*)
NCM	2008 ± 250(*ABC*)	931 ± 103(*ABC*)
OMd	521 ± 84(*A*)	–
OMv	449 ± 10(*A*)	–
pHVC	3539 ± 479(*ABC*)	1826 ± 350(*BC*)
Uva	1523 ± 132(*ABC*)	627 ± 48(*ABC*)
RA	1756 ± 83(*ABC*)	805 ± 122(*ABC*)
TeO	792 ± 50(*ABC*)	585 ± 118(*ABC*)
TnA	1045 ± 169(*ABC*)	998 ± 268(*ABC*)
VMH	935 ± 210(*ABC*)	914 ± 469(*B*)

Letters (A, B, or C) represent individual birds that showed binding in the indicated brain area (–, area not distinguishable from background binding).

TABLE 4 | Mean optical binding density (±SEM) of ^{125}I-OVTA and ^{125}I-LVA in brain regions with visible binding in rock dove.

Brain region	^{125}I-OVTA	^{125}I-LVA
A	662 ± 100(*BC*)	–
APH	983 ± 223(*ABC*)	765 ± 145(*BC*)
Bas	1534 ± 364(*ABC*)	–
CMM	801 ± 196(*ABC*)	743 ± 149(*ABC*)
DLP	1628 ± 73(*ABC*)	1228 ± 97(*BC*)
DMP	793 ± 46(*BC*)	–
E	671 ± 75(*ABC*)	–
H	540 ± 54(*BC*)	–
Hp	1084 ± 97(*BC*)	643 ± 117(*B*)
LHy	411 ± 57(*BC*)	–
LS	2692 ± 221(*ABC*)	936 ± 221(*ABC*)
M	612 ± 92(*BC*)	–
MSt	3613 ± 235(*C*)	890 ± 188(*C*)
MVL	1493 ± 145(*BC*)	483 ± 57(*B*)
NIM	1596 ± 156(*BC*)	1270 ± 397(*C*)
SGP	576 ± 117(*ABC*)	–
TeO	312 ± 36(*ABC*)	542 ± 108(*B*)

Letters (A, B, or C) represent individual birds that showed binding in the indicated brain area (–, area not distinguishable from background binding).

Distribution of ^{125}I-OVTA and ^{125}I-LVA Binding in House Sparrow

Table 2 provides a complete list of brain regions that presented with visible radioligand binding, and **Figure 1** shows representative autoradiograms of ^{125}I-OVTA and ^{125}I-LVA binding sites in house sparrows. ^{125}I-OVTA and ^{125}I-LVA binding was limited to relatively few brain regions, although labeling was widely dispersed across the rostral-caudal axis of the brain. Unlike in European starlings and rock doves, in which the highest level of binding was observed in portions of the septal complex, in house sparrows the densest ^{125}I-OVTA signal was observed in the Ad, a trend that was noted across all three individuals. Other areas showing dense ^{125}I-OVTA binding include the parahippocampal area (APH), the LS, and the para-high vocal center (pHVC). Binding in the LS and APH occurred across all three individuals, while pHVC labeling was observed in two out of three subjects. ^{125}I-LVA binding occurred only in sites that also showed ^{125}I-OVTA binding. In addition, the presence of ^{125}I-LVA binding was highly variable across individuals, such that ^{125}I-LVA signal occurred in some, but not all individuals for all regions except for the optic tectum (TeO).

Distribution of ^{125}I-OVTA and ^{125}I-LVA Binding in European Starling

Representative autoradiograms and a complete list of brain areas in the European starling with visible binding appear in **Figure 2** and **Table 3**, respectively. All three female starlings showed high levels of ^{125}I-OVTA binding in the LS, Ad, lateral arcopallium (Al), and pHVC, with the strongest signals occurring in the LS. More moderate ^{125}I-OVTA binding occurred across all three females in portions of the nidopallium, especially the caudal
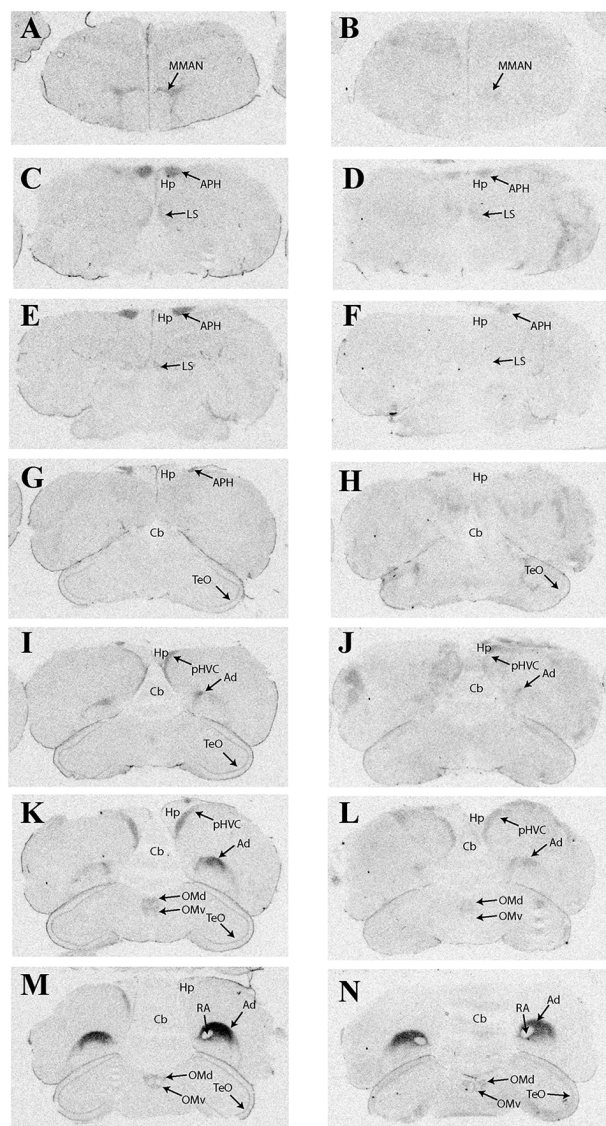
FIGURE 1 | Representative photomicrographs of ^{125}I-ornithine vasotocin analog (^{125}I-OVTA; **A,C,E,G,I,K,M**) or ^{125}I-linearized vasopressin antagonist (^{125}I-LVA; **B,D,F,H,J,L,N**) binding in the brain of a house sparrow (images correspond to individual "B" in **Table 2**).
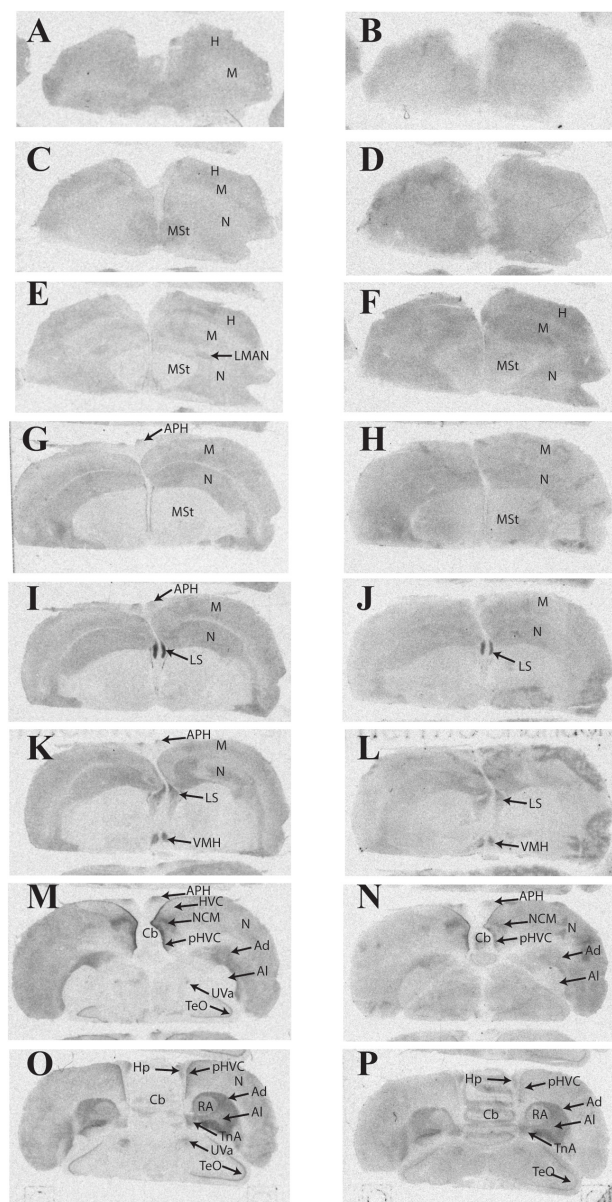

FIGURE 2 | Representative photomicrographs of ^{125}I-ornithine vasotocin analog (^{125}I-OVTA; **A,C,E,G,I,K,M,O**) or ^{125}I-linearized vasopressin antagonist (^{125}I-LVA; **B,D,F,H,J,L,N,P**) binding in the brain of a European starling (images correspond to individual "B" in **Table 3**).

medial region (NCM); the commissural septal nucleus (CoS); the uvaeform nucleus (Uva); robust nucleus of the arcopallium (RA); and the nucleus taeniae of the amygdala (TnA). As in house sparrows, ^{125}I-LVA binding almost exclusively overlapped with ^{125}I-OVTA binding sites, with the exception of the caudomedial mesopallium (CMM), which showed ^{125}I-LVA signal in one female. Similar to ^{125}I-OVTA, the highest density of ^{125}I-LVA binding sites occurred in the LS, Ad, Al, and pHVC, although binding did not appear in all three subjects for all of these regions. In comparison to ^{125}I-OVTA, the distribution of ^{125}I-LVA binding sites was more limited and, in all cases in which both radioligands bound to a region, ^{125}I-LVA binding density was lower.

Distribution of ^{125}I-OVTA and ^{125}I-LVA Binding in Rock Dove

Representative photomicrographs and a complete list of brain regions showing radioligand binding in rock doves appear in **Figure 3** and **Table 4**, respectively. As in European starlings, ^{125}I-OVTA binding was more broadly distributed and denser than ^{125}I-LVA binding. High levels of ^{125}I-OVTA labeling were noted in the medial striatum (MSt), basorostral pallial nucleus (Bas), dorsolateral nucleus of the posterior thalamus (DLP), hippocampus (Hp), LS ventrolateral nucleus of the mesopallium
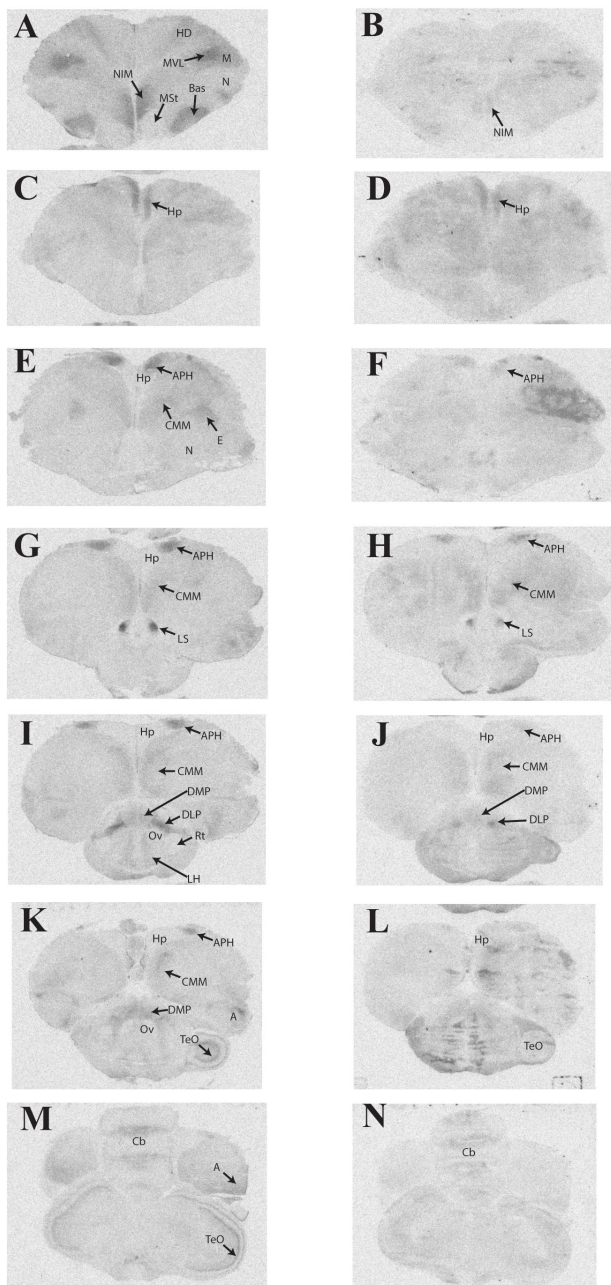
FIGURE 3 | Representative photomicrographs of [125]I-ornithine vasotocin analog ([125]I-OVTA; **A,C,E,G,I,K,M**) or [125]I-linearized vasopressin antagonist ([125]I-LVA; **B,D,F,H,J,L,N**) binding in the brain of a rock dove (images correspond to individual "B" in **Table 4**).

(MVL), and the intermediate medial nidopallium (NIM). Binding in several of these regions—MSt, Hp, MVL, and NIM—was only apparent in a subset of females. [125]I-LVA binding was highest in the NIM and DLP, though labeling in these regions appeared in fewer subjects compared to [125]I-OVTA. This trend—observing [125]I-LVA labeling in fewer subjects compared to [125]I-OVTA—was repeated across all brain areas except the CMM, LS, and

MSt. Regarding the latter, [125]I-LVA and [125]I-OVTA binding occurred in a distinct ring-like pattern in one female, whereas the other two subjects showed no observable labeling in MSt (**Figure 4**).

Competitive Binding Patterns in Lateral Septum of House Sparrows, European Starlings, and Rock Doves

The impacts of the Manning compound on binding patterns in the LS were strikingly consistent across all three species (house sparrows, **Figure 5**; European starlings, **Figure 6**; rock doves, **Figure 7**). In rock doves and European starlings, the competitor significantly reduced binding of both [125]I-OVTA [rock doves: $Z = 3.60$, $P = 0.0003$; European starlings: $t(8.67) = 7.77$, $P < 0.0001$] and [125]I-LVA [rock doves: $t(10.07) = 3.41$, $P = 0.007$; European starlings: $t(9.13) = 2.43$, $P = 0.04$]. Similar trends were observed in house sparrows, where the Manning compound induced significant and near significant reductions in binding for [125]I-OVTA ($Z = 3.24$, $P = 0.001$) and [125]I-LVA ($Z = 1.94$, $P = 0.05$), respectively. In all three species, [125]I-OVTA binding was significantly higher than [125]I-LVA in the absence of the Manning compound [rock doves: $t(16) = 5.62$, $P < 0.0001$; European starlings: $t(16) = 4.12$, $P = 0.0008$; house sparrows: $t(16) = 8.22$, $P < 0.0001$], but addition of the competitor eliminated this difference.

Competitive Binding Patterns in Dorsal Arcopallium of House Sparrows and European Starlings

Competitive binding patterns were similar across house sparrows (**Figure 8**) and European starlings (**Figure 9**). In the absence of the Manning compound in both sparrows and starlings, [125]I-OVTA binding in Ad was higher than [125]I-LVA, though this effect was significant for starlings ($Z = 3.53$, $P = 0.002$), but not sparrows ($Z = 1.59$, $P = 0.38$). Addition of the competitor significantly decreased [125]I-OVTA binding in both sparrows ($Z = 3.54$, $P = 0.002$) and starlings ($Z = 3.53$, $P = 0.002$). Addition of the Manning compound similarly reduced [125]I-LVA binding, though this effect was significant in European starlings ($Z = 3.53$, $P = 0.002$), but not house sparrows ($Z = 1.50$, $P = 0.44$). Although the Manning compound reduced binding of both radioligands, [125]I-LVA binding was significantly higher than [125]I-OVTA in the presence of the competitor, a trend that was observed in both sparrows ($Z = 2.65$, $P = 0.04$) and starlings ($Z = 3.53$, $P = 0.002$). In rock doves, two of three females showed low [125]I-OVTA binding in the arcopallium; [125]I-LVA binding was absent in this region. Thus, competitive binding patterns were not assessed in the arcopallium in rock doves.

DISCUSSION

The goals of our research were twofold: first, to establish neuroanatomical maps of NP receptors in three promising models for neuroecological examinations of collective behavior, and second, to examine the composition of NP receptor
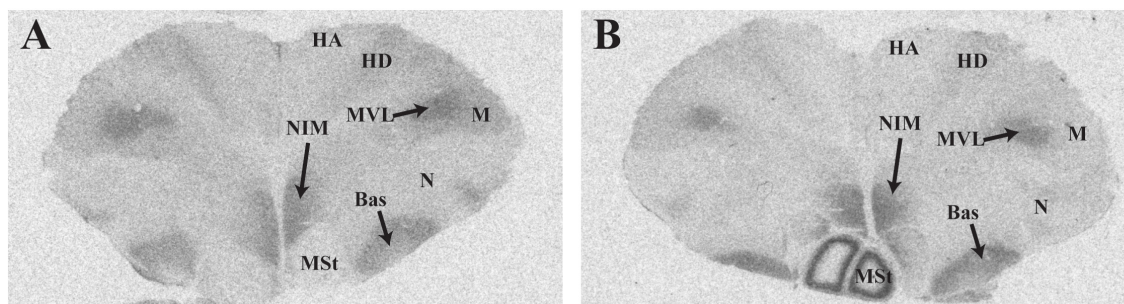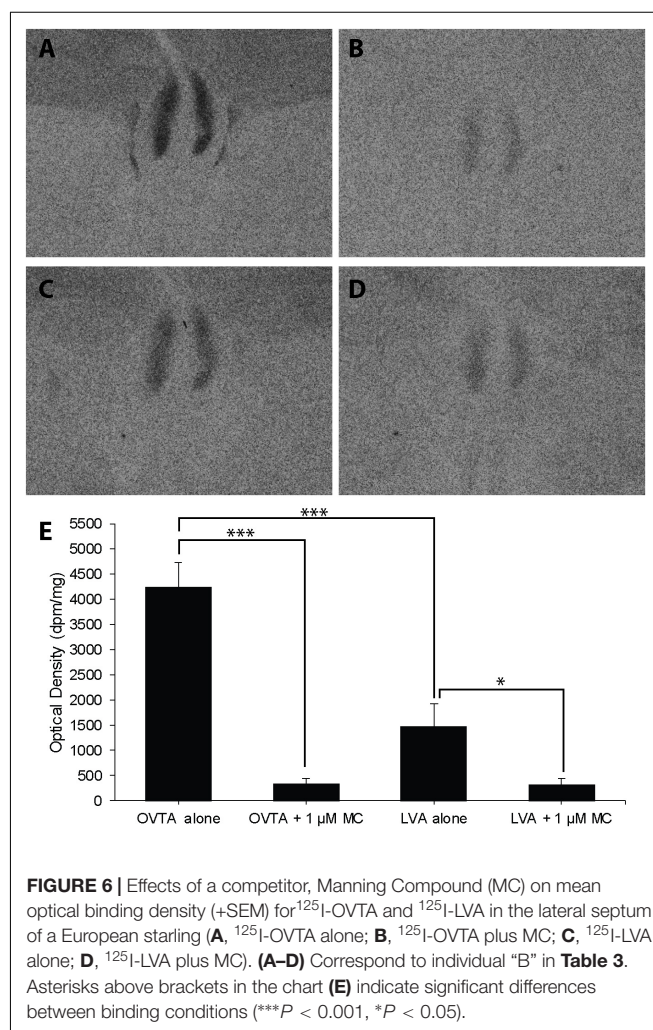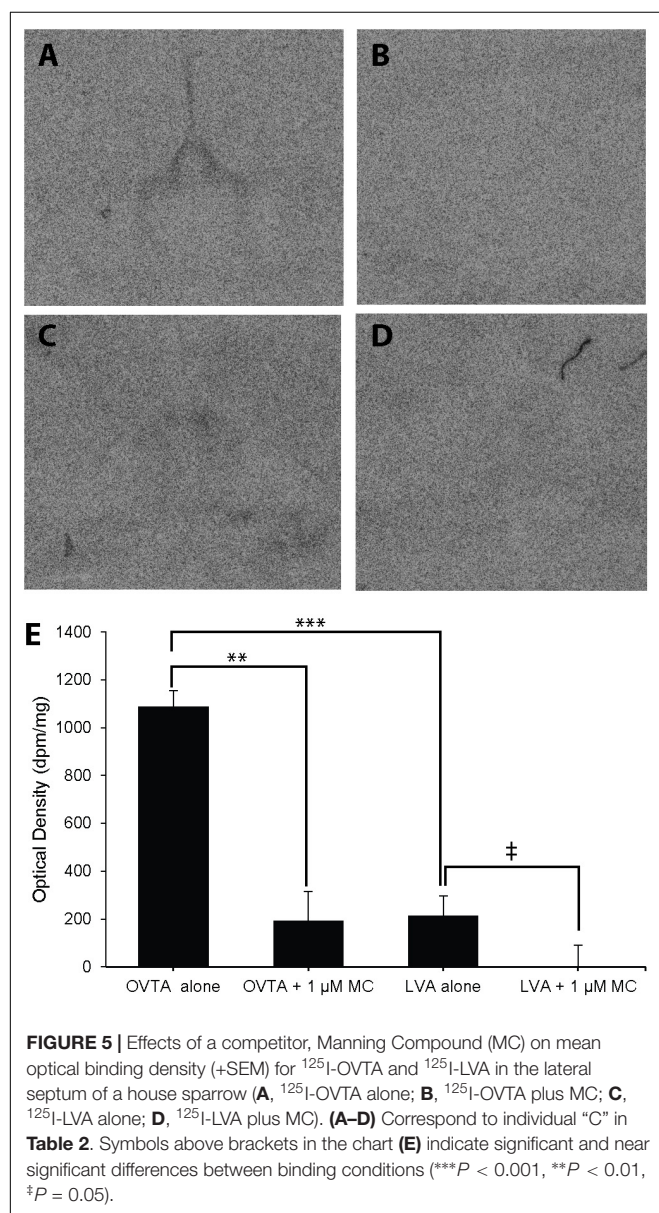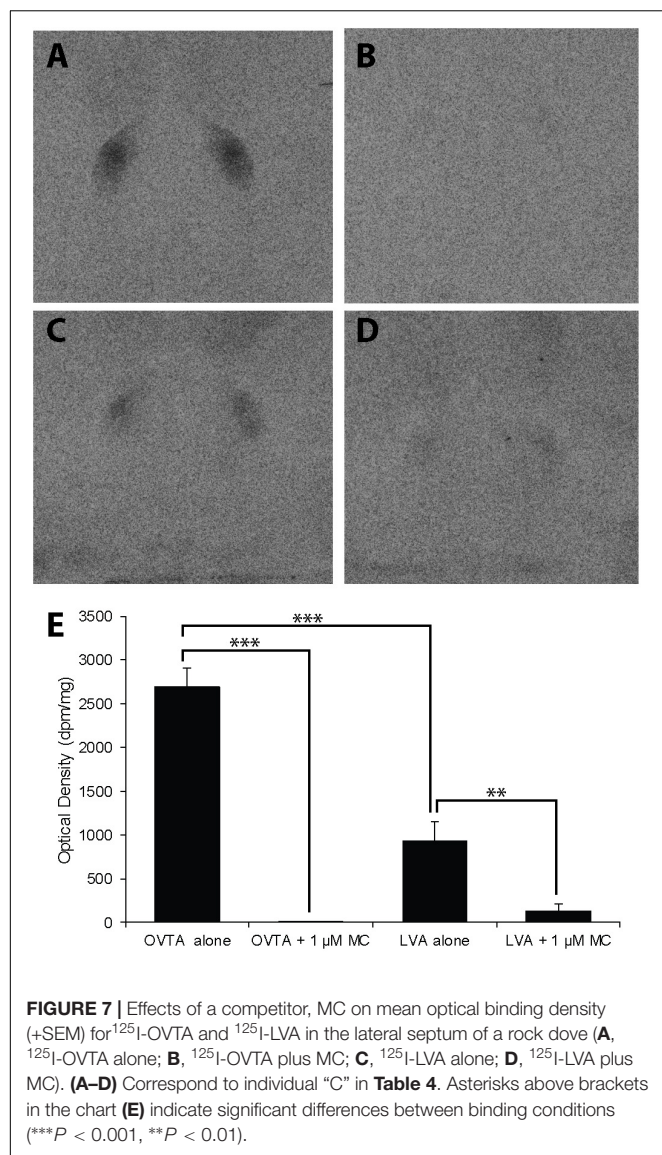
FIGURE 4 | Photomicrographs showing diverse binding patterns of ^{125}I-ornithine vasotocin analog (^{125}I-OVTA) in medial striatum of two female rock doves. For the bird represented by the right panel **(B)**, incubation with ^{125}I-linearized vasopressin antagonist (^{125}I-LVA) produced similar ring-like binding in the medial striatum, although the signal was less intense. Images in the left **(A)** and right **(B)** panels correspond with individuals "B" and "C," respectively, in **Table 4**. HA, apical hyperpallium; HD, densicellular hyperpallium.

FIGURE 5 | Effects of a competitor, Manning Compound (MC) on mean optical binding density (+SEM) for ^{125}I-OVTA and ^{125}I-LVA in the lateral septum of a house sparrow (**A**, ^{125}I-OVTA alone; **B**, ^{125}I-OVTA plus MC; **C**, ^{125}I-LVA alone; **D**, ^{125}I-LVA plus MC). **(A–D)** Correspond to individual "C" in **Table 2**. Symbols above brackets in the chart **(E)** indicate significant and near significant differences between binding conditions (***$P < 0.001$, **$P < 0.01$, ‡$P = 0.05$).

FIGURE 6 | Effects of a competitor, Manning Compound (MC) on mean optical binding density (+SEM) for ^{125}I-OVTA and ^{125}I-LVA in the lateral septum of a European starling (**A**, ^{125}I-OVTA alone; **B**, ^{125}I-OVTA plus MC; **C**, ^{125}I-LVA alone; **D**, ^{125}I-LVA plus MC). **(A–D)** Correspond to individual "B" in **Table 3**. Asterisks above brackets in the chart **(E)** indicate significant differences between binding conditions (***$P < 0.001$, *$P < 0.05$).

populations of these species using competitive binding. Our findings confirm our prediction that ^{125}I-LVA binding would be more limited than ^{125}I-OVTA binding and support the existence of multiple NP receptor types with overlapping

FIGURE 7 | Effects of a competitor, MC on mean optical binding density (+SEM) for[125]I-OVTA and [125]I-LVA in the lateral septum of a rock dove (**A**, [125]I-OVTA alone; **B**, [125]I-OVTA plus MC; **C**, [125]I-LVA alone; **D**, [125]I-LVA plus MC). **(A–D)** Correspond to individual "C" in **Table 4**. Asterisks above brackets in the chart **(E)** indicate significant differences between binding conditions (***$P < 0.001$, **$P < 0.01$).

FIGURE 8 | Effects of a competitor, MC on mean optical binding density (+SEM) for[125]I-OVTA and [125]I-LVA in the arcopallium of a house sparrow (**A**, [125]I-OVTA alone; **B**, [125]I-OVTA plus MC; **C**, [125]I-LVA alone; **D**, [125]I-LVA plus MC). **(A–D)** Correspond to individual "B" in **Table 2**. Asterisks above brackets in the chart **(E)** indicate significant differences between binding conditions (**$P < 0.01$, *$P < 0.05$).
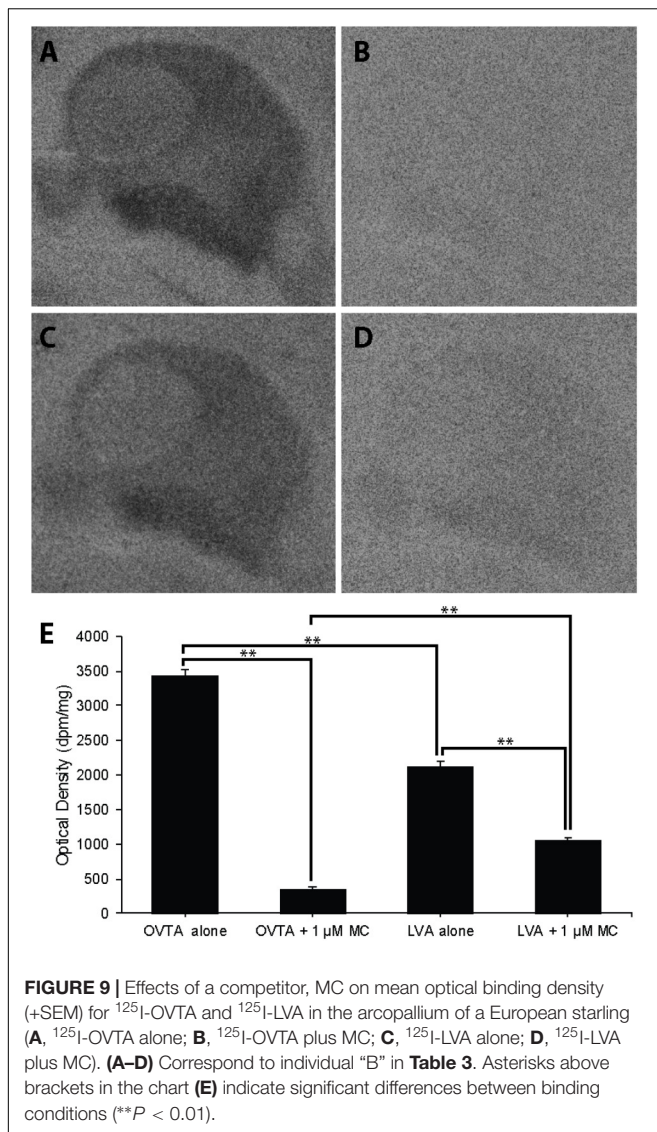
distributions. Below, we discuss binding patterns in rock doves, European starlings, and house sparrows in the context of NP receptor maps reported for other avian species; discuss the functional implications of binding in specific brain areas; and discuss the implications of our work for future neuroecological investigations of grouping behaviors. Although this study was not designed to provide a robust quantitative test of interspecies differences in NP receptor density or distribution, qualitative examination of our results highlights potentially valuable, novel lines of inquiry for understanding the neuroecological bases of collective behaviors, which we discuss further below.

Radioligand Binding in an Interspecies Context

Similar to reports in other avian species (Goodson et al., 2006; Leung et al., 2009), we found that [125]I-LVA binding was limited
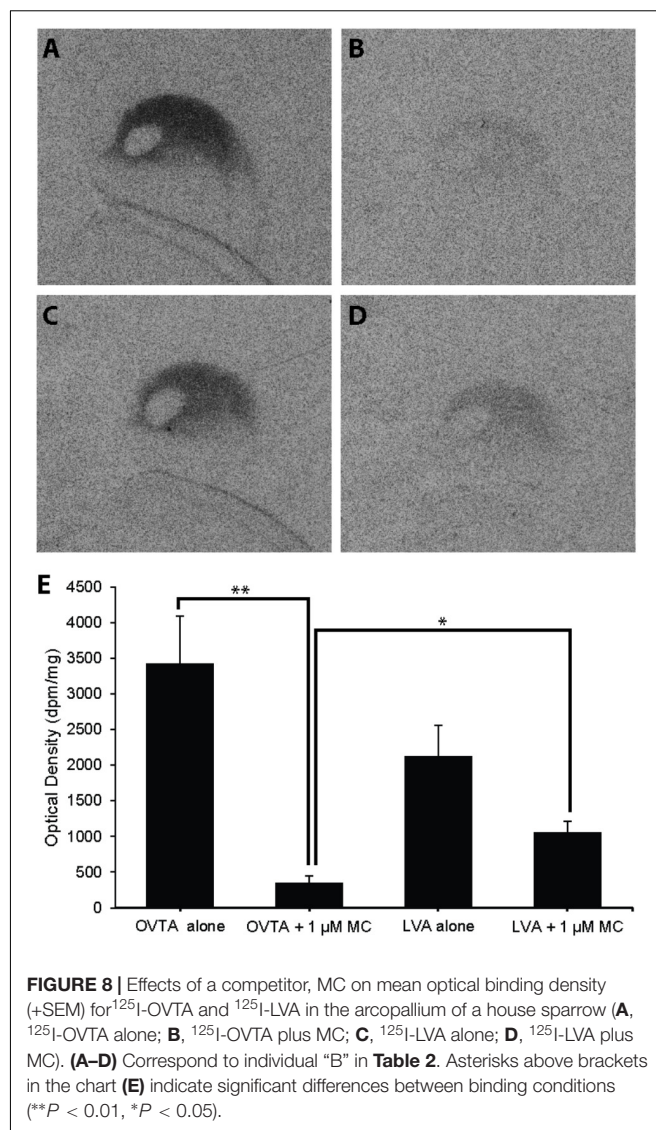
across all three species. Specifically, we found that [125]I-LVA signal appeared in fewer brain regions, in fewer individuals, and at lower densities when compared to [125]I-OVTA. In house sparrows and European starlings, the most pronounced [125]I-LVA binding occurred in portions of the arcopallium, while in rock doves, the highest level of [125]I-LVA signal appeared in the DLP and NIM, though only a subset of individuals showed binding in these regions. All three species showed [125]I-LVA binding in LS. These results replicate similar findings of limited [125]I-LVA binding, often restricted to LS, in other avian species. For example, among several flocking and territorial Estrildid finch species [melba finch (*Pytilia melba*), violet-eared waxbill (*Uraeginthus granatina*), Angolan blue waxbill (*Uraeginthus angolensis*), spice finch (*Lonchura punctulata*), and zebra finch (*Taeniopygia guttata*)], only the spice finch shows pronounced binding outside of the LS (Goodson et al., 2006). Similarly, in the white-throated sparrow (*Zonotrichia albicollis*), [125]I-LVA binding is restricted to the septal nuclei, Ad, and TeO (Leung et al.,

FIGURE 9 | Effects of a competitor, MC on mean optical binding density (+SEM) for [125]I-OVTA and [125]I-LVA in the arcopallium of a European starling (**A**, [125]I-OVTA alone; **B**, [125]I-OVTA plus MC; **C**, [125]I-LVA alone; **D**, [125]I-LVA plus MC). (**A–D**) Correspond to individual "B" in **Table 3**. Asterisks above brackets in the chart (**E**) indicate significant differences between binding conditions (**$P < 0.01$).

binding patterns in other songbird species. For example, in the white-throated sparrow and zebra finch, [125]I-OVTA binds to receptors in the LS, TnA, APH, and arcopallium (Leung et al., 2009), and in several species of emberizid sparrow (field sparrow (*Spizella pusilla*), dark-eyed junco (*Junco hyemalis*), song sparrow (*Melospiza melodia*), and eastern towhee (*Pipilo erythrophthalamus*), [125]I-OVTA binds to the LS and arcopallium (Wilson et al., 2016). As in European starlings and house sparrows, [125]I-OVTA binding in rock doves was high in LS, but the overall [125]I-OVTA binding pattern showed several distinctions in rock doves relative to the other two species. Specifically, in rock doves, high [125]I-OVTA binding appeared in Bas, DLP, Hp, MVL, and NIM, but not in the arcopallium. In addition, one rock dove showed a striking and, to our knowledge, previously unreported distribution of NP receptors in a ring-like pattern along the MSt's outer margins.

The distinct binding patterns in the brain of the rock dove, when compared to European starlings and house sparrows, may have a variety of underlying causes, including evolutionarily driven interspecies differences, differences in the season of specimen collection, or differences in natural history or life history stage. Regarding the first explanation, it is worth noting that European starlings and house sparrows are both songbirds and more evolutionarily related to one another than to rock doves, which is a Columbiforme (an order of birds that includes pigeons and doves; Johnston and Janiga, 1995). Because our study was not designed to elucidate interspecies differences in NP receptor maps, future work will be needed to examine the validity of these explanations. Approximately half of all extant avian species are *not* songbirds (Barker et al., 2004); however, thus far all studies examining the relationship between NPs and grouping behavior in birds have used songbird species. Comparisons across both songbird and non-songbird taxa are needed to augment our understanding of the neural mechanisms that underlie flocking, as well as the generality of these mechanisms across avian species.

Grouping Behavior and NP Receptors in the Lateral Septum

In both mammals and birds, the LS appears to play an important role in regulating intra- and interspecies differences in social behavior. For example, female meadow voles (*Microtus pennsylvanicus*), which form groups in winter, show variations in same-sex huddling that are associated with OTR expression in the LS (Beery and Zucker, 2010). Similarly, social (*Ctenomys sociabilis*) and solitary (*C. haigi*) species of rodents known as tuco–tucos show differences in OTR binding in LS (Beery et al., 2008a). In the zebra finch, NP receptors in the septal complex are associated with variations in group size preference (Goodson et al., 2009b). In addition, interspecies comparisons of estrildid finches show that [125]I-LVA and [125]I-OVTA binding in the caudal zone of the LS is higher in flocking versus territorial species, and infusions of V1aR and OTR antagonists directly into the zebra finch LS significantly decrease the duration of time that individuals spend near a large group of conspecifics (Goodson et al., 2006, 2009b; Kelly et al.,

2009). Although these avian taxa show similar [125]I-LVA binding patterns, they display varying degrees of grouping behavior, suggesting some degree of evolutionary conservation in brain-wide distribution for the receptor, or receptors, to which [125]I-LVA binds. However, variations in NP receptor distribution or density within specific brain regions may contribute to behavioral differences. For example, localized [125]I-LVA binding *within* septal areas has been associated with differences in grouping behavior among flocking and territorial avian species; similar findings have also been reported for [125]I-OVTA (Goodson et al., 2006, 2009b).

In contrast to [125]I-LVA, [125]I-OVTA binding was more intense and widely distributed in all three species. In house sparrows and European starlings, moderate to high levels of [125]I-OVTA binding occurred in the arcopallium, APH, septal areas, and pHVC. European starlings showed dense [125]I-OVTA binding in additional brain areas, including the NCM and TnA. The distribution of [125]I-OVTA binding in European starlings and house sparrows showed a number of similarities to [125]I-OVTA

2011). Intriguingly, variations in mesotocin innervation, but not NP receptor densities, in the LS are associated with different seasonal patterns of flocking behavior (i.e., flocking year-round versus winter flocking) across species of emberizid sparrows (Goodson et al., 2012; Wilson et al., 2016). These findings suggest that interspecies and seasonal variations in flocking may be differentially mediated by the NP systems, and highlight the importance of avoiding the assumption that a single mechanism governs apparently similar behavioral patterns. They also support consideration of brain areas other than the LS as potential mediators of seasonal variations in flocking.

Brain Areas With Unknown Contributions to Flocking: The Arcopallium and Sensory Pathways

Much focus has been placed on the LS in studies of NPs and their role in avian grouping behavior; however, several brain areas other than the LS also show dense expression of NP receptors. For example, using multiple songbird species, Leung et al. (2009) and Wilson et al. (2016) found moderate to high concentrations of NP receptors in the arcopallium and the caudal nidopallium, an area involved in auditory perception. Similarly, we found dense NP receptor expression in the arcopallium, particularly in the dorsal zone (in house sparrows and European starlings, but not rock doves), and in the caudomedial nidopallium (NCM; in European starlings). Little is known about the function of NP receptors in these brain regions, although there is reason to suspect that they may be involved in mediating social behavior. For example, Wilson et al. (2016) identified the rostral arcopallium as a potential "affiliation hot spot" in the avian brain because of its putative homology to the mammalian pallial amygdala, a region with well-established contributions to social behavior (Jarvis, 2009; Hanics et al., 2016). Furthermore, Wilson et al. (2016) found that seasonally flocking, but not non-flocking species of emberizid sparrows, show higher ^{125}I-OVTA binding in the rostral arcopallium during winter. In combination with our results, these findings implicate NPs in the arcopallium as potential mediators of seasonal variations in flocking, and support future investigations of this possibility in European starlings and house sparrows, but not in rock doves, which did not show robust radioligand binding in the arcopallium.

Although NP receptors have been previously identified in the NCM, it remains unknown whether NPs in this brain area mediate social interactions. In songbirds, the NCM is a key site for auditory processing and song control, as are several other brain regions, including the CMM, LMAN, MMAN, Uva, and RA (Foster and Bottjer, 1998). We observed NP receptor expression in all of these areas. Specifically, we found high binding density in the NCM and RA (in European starlings), and low binding density, or binding in only a subset of individuals per species, in the CMM (all three species), LMAN and Uva (in European starlings), and MMAN (in house sparrows). These regions are components of an interconnected song control system that governs song learning and maintenance (Foster

et al., 1997; Foster and Bottjer, 1998). Interestingly, we did not observe robust radioligand binding in Area X or the high vocal center (HVC), both of which constitute key sites in this network (Ziegler and Marler, 2008; Ellis and Riters, 2013). However, in house sparrows and European starlings, we observed dense NP receptor expression in the pHVC, a thin strip of cells that lines the medial edge of the HVC and lies within the margins of the NCM. Although the HVC and pHVC are neuroanatomically adjacent to each other, neural tracing studies show that the afferent and efferent projections of the pHVC are distinct from the HVC, suggesting that these two areas may be functionally distinct (Foster and Bottjer, 1998).

The social implications of NP expression in auditory and vocal brain regions remain almost wholly uninvestigated, although it is well established that NPs impact vocal behavior and learning across multiple taxa (e.g., fish: Goodson and Bass, 2000; mammals: Scattoni et al., 2008; Lukas and Wöhr, 2015; birds: Voorhuis et al., 1991; Maney et al., 1997; Goodson, 1998; Harding and Rowe, 2003; Goodson et al., 2009a; Baran et al., 2017). In the zebra finch—the most commonly used model for investigating neural control of singing behavior—pair bonding is correlated with the activation of NP receptor expression in auditory brain regions. Specifically, V1aR-like, but not OTR-like, mRNA expression is higher in the NCM and CMM of paired, relative to unpaired, females (Tomaszycki and Atchley, 2017). In combination with our results, such findings indicate that the contributions of NPs in the song control network to social grouping constitutes a fruitful potential line of research, particularly since vocalizations are likely a key driver of group formation and maintenance, at least in some avian species. For example, across different social contexts, house sparrows display distinct vocalizations, including the "flock call," which is most readily observed in winter groups and appears to contribute to group cohesion, and a repetitive "chirrup," which is used by both males and females to facilitate the formation of foraging groups (Elgar, 1986a,b; Anderson, 2006).

Interestingly, rock doves, but not European starlings or house sparrows, displayed high levels of binding in several brain areas that are involved in sensory pathways, including the MSt. The structural basis and functional implications of the ring-like binding pattern in the MSt of the rock doves are unclear. However, the avian MSt is known to be a heterogenous area that is composed of multiple cell types, with connectivity and neurochemical traits that differ on a mediolateral axis. Specifically, neural tracing studies implicate the medial MSt in viscerolimbic processes—which facilitate the translation of contextual stimuli into behavioral responses (Goodson and Kabelik, 2009; Kuenzel et al., 2011)—and the lateral MSt in somatosensory, visual, auditory, and motor function. We also found dense binding in additional brain nuclei—including the MVL, NIM, Bas, and DLP— that are interconnected by sensory pathways involved in transmitting visual, somatosensory, and auditory information (Atoji and Wild, 2012). Johnson and Young (2017) report that diverse taxa display NP receptors in sensory nuclei and posit that the distribution of receptors in these

areas reflects "dominant socio-sensory modalities" used by each species.

Evidence for Distinct Receptor Populations in the Avian Brain

Intraspecies comparisons of binding distributions, using multiple radioligands and binding competitors, are needed to identify heterogenous populations of NP receptors (Leung et al., 2009). However, studies that map and compare binding for both ^{125}I-LVA and ^{125}I-OVTA—perhaps the two most frequently used radioligands for NP receptor mapping in both mammalian and avian species—have only been conducted in two avian species: the white-throated sparrow and zebra finch (Leung et al., 2009). As described previously, Leung et al. (2009) found that ^{125}I-OVTA binding was more widespread than ^{125}I-LVA binding. However, the regions to which these radioligands bound were highly overlapping, leading the authors to conclude that their results could support either the presence of multiple NP receptor types, or a single receptor with differing levels of affinity for ^{125}I-LVA and ^{125}I-OVTA.

Although our findings regarding ^{125}I-LVA and ^{125}I-OVTA binding patterns are strikingly similar to Leung et al. (2009), the results from our competitive binding experiment suggest that the NP receptor populations in the LS versus the arcopallium may be composed of different receptor subtypes. Specifically, we found in sparrows and starlings that the Manning compound reduces ^{125}I-LVA binding to a greater extent in the arcopallium than in the LS. Thus, our results suggest that the Manning compound affects radioligand binding in a brain region-specific manner, even in different species. However, our results also support the interpretation that the NP receptor subtypes show some degree of promiscuous radioligand binding. The binding that remains in the presence of the Manning compound may be due to radioligand binding at a different receptor subtype, likely the OTR-like receptor, VT3. If the Manning compound is indeed binding selectively to the avian V1a-like receptor, as it does primarily in mammalian systems, then our results would indicate that both radioligands bind to the V1a-like receptor, but also exhibit some affinity for the OTR-like avian NP receptor. The pharmacological cross-talk of ^{125}I-LVA and ^{125}I-OVTA to OTR and V1aR is already an established phenomenon which has been demonstrated in primate brain tissue (Freeman et al., 2014), and distinguishing NP receptor subtypes in primates is an active area of ongoing research. Our results support the idea that similar *in vitro* pharmacological investigations are merited in avian models as well.

Implications for Studies Using Manning Compound as an NP Receptor Antagonist

Because of its potency as a V1aR antagonist, the Manning compound is frequently used to examine the function of specific NP receptor types in mammals, although Manning et al. (2012) caution that this compound is also a potent *in vitro* OTR antagonist and "fairly potent" *in vivo* OTR antagonist. Nonetheless, the Manning compound has also become widely

used to identify the contributions of V1a-like receptors to a variety of avian social behaviors, including aggression, social attachment and affiliation, song learning, and pair maintenance behaviors (Goodson et al., 2004; Baran et al., 2016a,b, 2017).

Although the Manning compound is now a commonly used tool for determining the contributions of putative V1a-like receptors to the mediation of avian social behavior, the selectiveness of the Manning compound for specific NP receptors in the avian brain remains unknown. To the best of our knowledge, our study is the first to directly examine how the Manning compound impacts NP receptor binding in specific avian brain regions. We found that the Manning compound displaced ^{125}I-OVTA more readily than ^{125}I-LVA, which calls into question whether ^{125}I-OVTA is labeling OTR-like NP receptors in avian brains. This finding could also indicate that the Manning compound is not specifically targeting V1a-like receptors in avian models, which merits caution for studies using it to examine V1a-like receptor functions. We suggest conservative interpretations of ^{125}I-OVTA and ^{125}I-LVA binding distributions, until more extensive pharmacological studies are completed with avian NP receptors. We also suggest that future avian behavioral studies using the Manning compound as an antagonist should include two treatment groups: one that combines the Manning compound with vasotocin, and one that combines it with mesotocin. This experimental paradigm has been used previously to determine if the Manning compound selectively reverses the effects of vasotocin or mesotocin on avian behavior, and to provide an added test of the hypothesis that a specific receptor subtype is predominantly involved in behavioral mediation (Goodson et al., 2004).

Conclusion: The Value of Developing Avian Models for Social Neuroecology

Ecological conditions markedly influence social grouping across a diversity of species, but only a few studies have examined how ecological and neurobiological factors interact to mediate this behavior. Work with Amargosa pupfish (*Cyprinodon nevadensis*) indicate that social interactions are sensitive to the physical environment, and that NPs in the brain likely facilitate this ecological sensitivity. Specifically, exposure to high salinity in this species facilitates a reduction in the number of magnocellular vasotocin (VT) neurons in the preoptic area, and VT neuronal phenotypes, as well as aggression, vary with temperature regime (Lema, 2006). In addition, female meadow voles, which only form groups during winter, display same-sex affiliative behavior and group-size preference that varies with day length, temperature, and food restriction, as well as OTR binding that varies with day length (Beery et al., 2008b; Beery and Zucker, 2010; Ondrasek et al., 2015).

Similarly, European starlings, house sparrows, and rock doves display grouping behaviors that vary across ecological contexts. Notably, the behavioral profiles of each species are different in key ways that make each species advantageous for investigating particular questions about the neuroecology of collective behaviors. For example, starlings show striking seasonal patterns in flocking, such that they display high levels of aggression toward conspecifics during the breeding

period, but aggregate into highly coordinated flocks that may number in the millions during the winter months (Cabe, 1993; Goodenough et al., 2017). This observation raises the question of how seasonal environmental factors, particularly day length, influence neurochemical mediation of social coordination among individual starlings. House sparrows show less striking seasonal changes in flocking than starlings; however, throughout the year, they form temporary foraging flocks that vary in size according to the divisibility of a food source, the perceived risk of predation, and distance from cover (Elgar, 1986a,b, 1987), suggesting that food availability, in relation to other environmental factors, may impact the neural mechanisms underlying flock formation. Lastly, unlike European starlings and house sparrows, rock doves are commonly found breeding in large colonies in which individuals show some degree of behavioral coordination (e.g., flushing from their nests together in response to a predator). The size, composition, and location of such colonies vary with several environmental factors, including the availability of food and nest sites (Johnston and Janiga, 1995). Thus, rock doves present an opportunity to investigate the neural mechanisms underlying colonial breeding, how these mechanisms are influenced by ecological variations, and how neuroecological regulation of colonial breeding impacts reproductive success.

To conclude, the three species examined here serve as ideal models for neuroecological research for multiple reasons: they inhabit a wide range of environments, show grouping behaviors that vary across ecological contexts, and display NP receptors in brain regions that may play a role in avian flocking. In addition, because several aspects of the NP systems are evolutionarily conserved across vertebrate taxa (Gimpl and Fahrenholz, 2001; Goodson, 2005, 2013), discoveries made using these species may guide the development of hypotheses and predictions for subsequent investigations across a much wider array of taxa.

AUTHOR CONTRIBUTIONS

NO and RC completed field collections of specimens. NO and SF were jointly responsible for overall experimental design and execution of the autoradiography assay. NO scored films and wrote the manuscript in conjunction with SF, with editorial suggestions from KB and RC.

FUNDING

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fnins.2018.00713/full#supplementary-material

REFERENCES

Acher, R., Chauvet, J., and Chauvet, M. T. (1995). Man and the chimaera. Selective versus neutral oxytocin evolution. *Adv. Exp. Med. Biol.* 395, 615–627.

Alexander, R. D. (1974). The evolution of social behavior. *Ann. Rev. Ecol. Syst.* 5, 325–383. doi: 10.1146/annurev.es.05.110174.001545

Anderson, T. R. (2006). *Biology of the Ubiquitous House sparrow: From Genes to Populations.* New York, NY: Oxford University Press. doi: 10.1093/acprof:oso/9780195304114.001.0001

Atoji, Y., and Wild, J. M. (2012). Afferent and efferent projections of the mesopallium in the pigeon (*Columba livia*). *J. Comp. Neurol.* 520, 717–741. doi: 10.1002/cne.22763

Baran, N. M., Peck, S. C., Kim, T. H., Goldstein, M. H., and Adkins-Regan, E. (2017). Early life manipulations of vasopressin-family peptides alter vocal learning. *Proc. R. Soc. B* 284:20171114. doi: 10.1098/rspb.2017.1114

Baran, N. M., Sklar, N. C., and Adkins-Regan, E. (2016a). Developmental effects of vasotocin and nonapeptide receptors on early social attachment and affiliative behavior in the zebra finch. *Horm. Behav.* 78, 20–31. doi: 10.1016/j.yhbeh.2015.10.005

Baran, N. M., Tomaszycki, M. L., and Adkins-Regan, E. (2016b). Early life manipulations of the nonapeptide system alter pair maintenance behaviors and neural activity in adult male zebra finches. *Front. Behav. Neurosci.* 10:58. doi: 10.3389/fnbeh.2016.00058

Barker, F. K., Cibois, A., Schikler, P., Feinstein, J., and Cracraft, J. (2004). Phylogeny and diversification of the largest avian radiation. *Proc. Natl. Acad. Sci. U.S.A.* 101, 11040–11045. doi: 10.1073/pnas.0401892101

Beery, A. K., Kamal, Y., Sobrero, R., and Hayes, L. D. (2016). "Comparative neurobiology and genetics of mammalian social behavior," in *Sociobiology of Caviomorph Rodents*, eds L. A. Ebensperger and L. D. Hayes (Hoboken, NJ: John Wiley & Sons, Ltd), 59–90.

Beery, A. K., Lacey, E. A., and Francis, D. D. (2008a). Oxytocin and vasopressin receptor distributions in a solitary and a social species of tuco-tuco (*Ctenomys haigi* and *Ctenomys sociabilis*). *J. Comp. Neurol.* 507, 1847–1859. doi: 10.1002/cne.21638

Beery, A. K., Loo, T. J., and Zucker, I. (2008b). Day length and estradiol affect same-sex affiliative behavior in the female meadow vole. *Horm. Behav.* 54, 153–159. doi: 10.1016/j.yhbeh.2008.02.007

Beery, A. K., and Zucker, I. (2010). Oxytocin and same-sex social behavior in female meadow voles. *Neuroscience* 169, 665–673. doi: 10.1016/j.neuroscience.2010.05.023

Cabe, P. R. (1993). "European starling (*Sturnus vulgaris*), version 2.0," in *The Birds of North America*, eds A. F. Poole and F. B. Gill (Ithaca, NY: Cornell Lab of Ornithology), doi: 10.2173/bna.48

Chini, B., Mouillac, B., Ala, Y., Balestre, M. N., Trumpp-Kallmeyer, S., Hoflack, J., et al. (1995). Tyr115 is the key residue for determining agonist selectivity in the V1a vasopressin receptor. *EMBO J.* 14, 2176–2182. doi: 10.1002/j.1460-2075.1995.tb07211.x

Chini, B., Mouillac, B., Balestre, M. N., Trumpp-Kallmeyer, S., Hoflack, J., Hibert, M., et al. (1996). Two aromatic residues regulate the response of the human oxytocin receptor to the partial agonist arginine vasopressin. *FEBS Lett.* 397, 201–206. doi: 10.1016/S0014-5793(96)01135-0

Clergeau, P., Savard, J.-P. L., Mennechez, G., and Falardeau, G. (1998). Bird abundance and diversity along an urban-rural gradient: a comparative study between two cities on different continents. *Condor* 100, 413–425. doi: 10.2307/1369707

De Groof, G., George, I., Touj, S., Stacho, M., Jonckers, E., Cousillas, H., et al. (2016). A three-dimensional digital atlas of the starling brain. *Brain Struct. Funct.* 221, 1899–1909. doi: 10.1007/s00429-015-1011-1

Elgar, M. A. (1986a). House sparrows establish foraging flocks by giving chirrup calls if the resources are divisible. *Anim. Behav.* 34, 169–174. doi: 10.1016/0003-3472(86)90020-5

Elgar, M. A. (1986b). The establishment of foraging flocks in house sparrows: Risk of predation and daily temperature. *Behav. Ecol. Sociobiol.* 19, 433–438. doi: 10.1007/BF00300546

Elgar, M. A. (1987). Food intake rate and resource availability: Flocking decisions in house sparrows. *Anim. Behav.* 35, 1168–1176. doi: 10.1016/S0003-3472(87)80173-2

Ellis, J. M. S., and Riters, L. V. (2013). Patterns of FOS protein induction in singing female starlings. *Behav. Brain Res.* 237, 148–156. doi: 10.1016/j.bbr.2012.09.004

Emlen, S. T. (1982). The evolution of helping. I. An ecological constraints model. *Am. Nat.* 119, 29–39. doi: 10.1086/283888

Foster, E. F., and Bottjer, S. W. (1998). Axonal connections of the high vocal center and surrounding cortical regions in juvenile and adult male zebra finches. *J. Comp. Neurol.* 397, 118–138. doi: 10.1002/(SICI)1096-9861(19980720)397:1<118::AID-CNE9>3.0.CO;2-3

Foster, E. F., Mehta, R. P., and Bottjer, S. W. (1997). Axonal connections of the medial magnocellular nucleus of the anterior neostriatum in zebra finches. *J. Comp. Neurol.* 382, 364–381. doi: 10.1002/cne.903820305

Freeman, S. M., Inoue, K., Smith, A. L., Goodman, M. M., and Young, L. J. (2014). The neuroanatomical distribution of oxytocin receptor binding and mRNA in the male rhesus macaque (*Macaca mulatta*). *Psychoneuroendocrinology* 45, 128–141. doi: 10.1016/j.psyneuen.2014.03.023

Gimpl, G., and Fahrenholz, F. (2001). The oxytocin receptor system: structure, function, and regulation. *Physiol. Rev.* 81, 629–683. doi: 10.1152/physrev.2001.81.2.629

Goodenough, A. E., Little, N., Carpenter, W. S., and Hart, A. G. (2017). Birds of a feather flock together: insights into starling murmuration behaviour revealed using citizen science. *PLoS One* 12:e0179277. doi: 10.1371/journal.pone.0179277

Goodson, J. L. (1998). Territorial aggression and dawn song are modulated by septal vasotocin and vasoactive intestinal polypeptide in male field sparrows (*Spizella pusilla*). *Horm. Behav.* 34, 67–77. doi: 10.1006/hbeh.1998.1467

Goodson, J. L. (2005). The vertebrate social behavior network: evolutionary themes and variations. *Horm. Behav.* 48, 11–22. doi: 10.1016/j.yhbeh.2005.02.003

Goodson, J. L. (2013). Deconstructing sociality, social evolution and relevant nonapeptide functions. *Psychoneuroendocrinology* 38, 465–478. doi: 10.1016/j.psyneuen.2012.12.005

Goodson, J. L., and Bass, A. H. (2000). Forebrain peptides modulate sexually polymorphic vocal circuitry. *Nature* 403, 769–772. doi: 10.1038/35001581

Goodson, J. L., Evans, A. K., and Wang, Y. (2006). Neuropeptide binding reflects convergent and divergent evolution in species-typical group sizes. *Horm. Behav.* 50, 223–236. doi: 10.1016/j.yhbeh.2006.03.005

Goodson, J. L., and Kabelik, D. (2009). Dynamic limbic networks and social diversity in vertebrates: from neural context to neuromodulatory patterning. *Front. Neuroendocrinol.* 30:429–441. doi: 10.1016/j.yfrne.2009.05.007

Goodson, J. L., Lindberg, L., and Johnson, P. (2004). Effects of central vasotocin and mesotocin manipulations on social behavior in male and female zebra finches. *Horm. Behav.* 45, 136–143. doi: 10.1016/j.yhbeh.2003.08.006

Goodson, J. L., Rinaldi, J., and Kelly, A. M. (2009a). Vasotocin neurons in the bed nucleus of the stria terminalis preferentially process social information and exhibit properties that dichotomize courting and non-courting phenotypes. *Horm. Behav.* 55, 197–202. doi: 10.1016/j.yhbeh.2008.10.007

Goodson, J. L., Schrock, S. E., Klatt, J. D., Kabelik, D., and Kingsbury, M. A. (2009b). Mesotocin and nonapeptide receptors promote estrildid flocking behavior. *Science* 325, 862–866. doi: 10.1126/science.1174929

Goodson, J. L., Wilson, L. C., and Schrock, S. E. (2012). To flock or fight: neurochemical signatures of divergent life histories in sparrows. *Proc. Natl. Acad. Sci. U.S.A.* 109, 10685–10692. doi: 10.1073/pnas.1203394109

Gubrij, K. I., Chaturvedi, C. M., Ali, N., Cornett, L. E., Kirby, J. D., Wilkerson, J., et al. (2005). Molecular cloning of an oxytocin-like receptor expressed in the chicken shell gland. *Comp. Biochem. Physiol. B Biochem. Mol. Biol.* 142, 37–45. doi: 10.1016/j.cbpc.2005.05.011

Guoynes, C. D., Simmons, T. C., Downing, G. M., Jacob, S., Solomon, M., and Bales, K. L. (2018). Chronic intranasal oxytocin has dose-dependent effects on central oxytocin and vasopressin systems in prairie voles (*Microtus ochrogaster*). *Neuroscience* 369, 292–302. doi: 10.1016/j.neuroscience.2017.11.037

Hanics, J., Teleki, G., Alpár, A., Székely, A. D., and Csillag, A. (2016). Multiple amygdaloid divisions of arcopallium send convergent projections to the nucleus accumbens and neighboring subpallial amygdala regions in the domestic chicken: a selective pathway tracing and reconstruction study. *Brain Struct. Funct.* 222, 301–315. doi: 10.1007/s00429-016-1219-8

Harding, C. F., and Rowe, S. A. (2003). Vasotocin treatment inhibits courtship in male zebra finches; concomitant androgen treatment inhibits this effect. *Horm. Behav.* 44, 413–418. doi: 10.1016/j.yhbeh.2003.06.007

Hartman, S., Freeman, S. M., Bales, K. L., and Belsky, J. (2018). Prenatal stress as a risk—and an opportunity—factor. *Psychol. Sci.* 29, 572–580. doi: 10.1177/0956797617739983

Helm, B., Piersma, T., and van der Jeugd, H. (2006). Sociable schedules: interplay between avian seasonal and social behaviour. *Anim. Behav.* 72, 245–262. doi: 10.1016/j.anbehav.2005.12.007

Jarvis, E. D. (2009). "Evolution of the pallium in birds and reptiles," in *Encyclopedia of Neuroscience*, eds M. Binder, N. Hirokawa, and U. Windhorst (Berlin: Springer), 1390–1400. doi: 10.1007/978-3-540-29678-2_3165

Jarvis, E. D., Gunturkun, O., Bruce, L., Csillag, A., Karten, H., Kuenzel, W., et al. (2005). Avian brains and a new understanding of vertebrate brain evolution. *Nat. Rev. Neurosci.* 6, 151–159. doi: 10.1038/nrn1606

Johnson, Z. V., and Young, L. J. (2017). Oxytocin and vasopressin neural networks: implications for social behavioral diversity and translational neuroscience. *Neurosci. Biobehav. Rev.* 76, 87–98. doi: 10.1016/j.neubiorev.2017.01.034

Johnston, R. F., and Janiga, M. (1995). *Feral pigeons.* New York, NY: Oxford University Press.

Karten, H. J., and Hodos, W. (1967). *A Stereotaxic Atlas of the Brain of the Pigeon (Columba livia).* Baltimore: Johns Hopkins Press.

Kelly, A. M., Kingsbury, M. A., Hoffbuhr, K., Schrock, S. E., Waxman, B., Kabelik, D., et al. (2011). Vasotocin neurons and septal V1a-like receptors potently modulate songbird flocking and responses to novelty. *Horm. Behav.* 60, 12–21. doi: 10.1016/j.yhbeh.2011.01.012

Kruszynski, M., Lammek, B., Manning, M., Seto, J., Haldar, J., and Sawyer, W. H. (1980). [1-(.beta.-mercapto-.beta.,.beta.-cyclopentamethylenepropionic acid),2-(O-methyl)tyrosine]arginine-vasopressin and [1-(.beta.-mercapto-.beta.,.beta.-cyclopentamethylenepropionic acid)]arginine-vasopressin, two highly potent antagonists of the vasopressor response to arginine-vasopressin. *J. Med. Chem.* 23, 364–368. doi: 10.1021/jm00178a003

Kuenzel, W. J., Medina, L., Csillag, A., Perkel, D. J., and Reiner, A. (2011). The avian subpallium: new insights into structural and functional subdivisions occupying the lateral subpallial wall and their embryological origins. *Brain. Res.* 1424, 67–101. doi: 10.1016/j.brainres.2011.09.037

Lema, S. C. (2006). Population divergence in plasticity of the AVT system and its association with aggressive behaviors in a Death Valley pupfish. *Horm. Behav.* 50, 183–193. doi: 10.1016/j.yhbeh.2006.02.010

Leung, C. H., Abebe, D. F., Earp, S. E., Goode, C. T., Grozhik, A. V., Mididoddi, P., et al. (2011). Neural distribution of vasotocin receptor mRNA in two species of songbird. *Endocrinology* 152, 4865–4881. doi: 10.1210/en.2011-1394

Leung, C. H., Goode, C. T., Young, L. J., and Maney, D. L. (2009). Neural distribution of nonapeptide binding sites in two species of songbird. *J. Comp. Neurol.* 513, 197–208. doi: 10.1002/cne.21947

Lukas, M., and Wöhr, M. (2015). Endogenous vasopressin, innate anxiety, and the emission of pro-social 50-kHz ultrasonic vocalizations during social play behavior in juvenile rats. *Psychoneuroendocrinology* 56, 35–44. doi: 10.1016/j.psyneuen.2015.03.005

Maney, D. L., Goode, C. T., and Wingfield, J. C. (1997). Intraventricular infusion of arginine vasotocin induces singing in a female songbird. *J. Neuroendocrinol.* 9, 487–491. doi: 10.1046/j.1365-2826.1997.00635.x

Manning, M., Misicka, A., Olma, A., Bankowski, K., Stoev, S., Chini, B., et al. (2012). Oxytocin and vasopressin agonists and antagonists as research tools and

potential therapeutics. *J. Neuroendocrinol.* 24, 609–628. doi: 10.1111/j.1365-2826.2012.02303.x

Nixdorf-Bergweiler, B. E., and Bischof, H.-J. (2007). *A Stereotaxic Atlas of the Brain of the Zebra Finch, Taeniopygia Guttata, with Special Emphasis on Telencephalic Visual and Song System Nuclei in Transverse and Sagittal Sections.* Bethesda, MD: National Library of Medicine.

Ocampo Daza, D., Lewicka, M., and Larhammar, D. (2012). The oxytocin/vasopressin receptor family has at least five members in the gnathostome lineage, including two distinct V2 subtypes. *Gen. Comp. Endocrinol.* 175, 135–143. doi: 10.1016/j.ygcen.2011.10.011

Ondrasek, N. R., Wade, A., Burkhard, T., Hsu, K., Nguyen, T., Post, J., et al. (2015). Environmental modulation of same-sex affiliative behavior in female meadow voles (*Microtus pennsylvanicus*). *Physiol. Behav.* 140, 118–126. doi: 10.1016/j.physbeh.2014.12.021

Parrish, J. K., and Edelstein-Keshet, L. (1999). Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science* 284, 99–101. doi: 10.1126/science.284.5411.99

Perkeybile, A. M., Delaney-Busch, N., Hartman, S., Grimm, K. J., and Bales, K. L. (2015). Intergenerational transmission of alloparental behavior and oxytocin and vasopressin receptor distribution in the prairie vole. *Front. Behav. Neurosci.* 9:191. doi: 10.3389/fnbeh.2015.00191

Reiner, A., Perkel, D. J., Mello, C. V., and Jarvis, E. D. (2004). Songbirds and the revised avian brain nomenclature. *Ann. N. Y. Acad. Sci.* 1016, 77–108. doi: 10.1196/annals.1298.013

Ruxton, G. D., and Beauchamp, G. (2008). Time for some a priori thinking about post hoc testing. *Behav. Ecol.* 19, 690–693. doi: 10.1093/beheco/arn020

Scattoni, M. L., McFarlane, H. G., Zhodzishsky, V., Caldwell, H. K., Young, W. S., Ricceri, L., et al. (2008). Reduced ultrasonic vocalizations in vasopressin 1b knockout mice. *Behav. Brain Res.* 187, 371–378. doi: 10.1016/j.bbr.2007.09.034

Sherry, D. F. (2006). Neuroecology. *Ann. Rev. Psychol.* 57, 167–197. doi: 10.1146/annurev.psych.56.091103.070324

Solomon, N. G. (2003). A reexamination of factors influencing philopatry in rodents. *J. Mammal.* 84, 1182–1197. doi: 10.1644/BLe-013

Tomaszycki, M. L., and Atchley, D. (2017). Pairing increases activation of V1aR, but not OTR, in auditory regions of zebra finches: the importance of signal modality in nonapeptide-social behavior relationships. *Integr. Comp. Biol.* 57, 878–890. doi: 10.1093/icb/icx043

Voorhuis, T. A. M., De Kloet, E. R., and De Wied, D. (1991). Effect of a vasotocin analog on singing behavior in the canary. *Horm. Behav.* 25, 549–559. doi: 10.1016/0018-506X(91)90020-I

Wilson, L. C., Goodson, J. L., and Kingsbury, M. A. (2016). Seasonal variation in group size is related to seasonal variation in neuropeptide receptor density. *Brain Behav. Evol.* 88, 111–126. doi: 10.1159/000448372

Ziegler, H. P., and Marler, P. (2008). *Neuroscience of Birdsong.* Cambridge: Cambridge University Press.

Zimmer, R. K., and Derby, C. D. (2011). Neuroecology and the need for broader synthesis. *Integr. Comp. Biol.* 51, 751–755. doi: 10.1093/icb/icr070

Advantages of publishing in Frontiers

OPEN ACCESS
Articles are free to read for greatest visibility and readership

FAST PUBLICATION
Around 90 days from submission to decision

HIGH QUALITY PEER-REVIEW
Rigorous, collaborative, and constructive peer-review

TRANSPARENT PEER-REVIEW
Editors and reviewers acknowledged by name on published articles

Frontiers
Avenue du Tribunal-Fédéral 34
1005 Lausanne | Switzerland

Visit us: www.frontiersin.org
Contact us: info@frontiersin.org | +41 21 510 17 00

REPRODUCIBILITY OF RESEARCH
Support open data and methods to enhance research reproducibility

DIGITAL PUBLISHING
Articles designed for optimal readership across devices

FOLLOW US
@frontiersin

IMPACT METRICS
Advanced article metrics track visibility across digital media

EXTENSIVE PROMOTION
Marketing and promotion of impactful research

LOOP RESEARCH NETWORK
Our network increases your article's readership