frontiers
RESEARCH TOPICS

BIOINSPIRED SOLUTIONS TO THE CHALLENGES OF CHEMICAL SENSING

Topic Editors
Ramon Huerta and Thomas Nowotny

frontiers in
NEUROENGINEERING

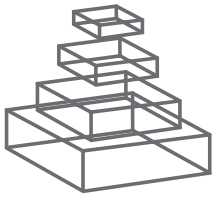

frontiers

Cover image provided by Ibbl sarl, Lausanne CH

ISSN 1664-8714
ISBN 978-2-88919-095-9
DOI 10.3389/978-2-88919-095-9

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

BIOINSPIRED SOLUTIONS TO THE CHALLENGES OF CHEMICAL SENSING

Topic Editors:
Ramon Huerta, University of California, San Diego, USA
Thomas Nowotny, University of Sussex, UK

Chemical sensing is likely the most primordial sensory modality that emerged in the evolution of life. Without chemical sensing life on earth would probably not exist. It is used for detecting nutrients, avoiding threats, finding mating partners and various forms of communication and social interaction between animals.

The advent of artificial sensors has created a myriad of problems in the areas of chemical detection and identification with applications in food quality and pollution control, chemical threat detection, health monitoring, robot control and even odor and taste synthesis. Efficient algorithms are needed to address the many challenges of chemical sensing in these areas, including (but not limited to) sensitivity levels, sensor drift, concentration invariance of analyte identity and complex mixtures. Defining and improving analysis methods for artificial chemical sensing remains an active research area in engineering and machine learning alike.

In the course of evolution animals, bacteria and plants have developed sophisticated methods and algorithms for solving difficult problems in chemical sensing very efficiently. Complex signaling pathways inside single cells can trigger movement toward the source of a nutrient. Complex networks of neurons appear to be able to compute odor types and the distance to a source in turbulent flows. These networks of neurons use a combination of temporal coding, layered structures, simple Hebbian learning rules, reinforcement learning and inhibition to quickly learn about chemical stimuli that are critical for their survival. Olfaction is a vibrant filed of research because recent technological advances allow monitoring and manipulating brain areas inaccessible in the past thus allowing for rapid progress. This is particularly relevant because to this date the best solutions to many general chemical sensing problems are still found in animals rather than artificial devices. Many lessons may yet have to be learned from biological systems to solve the complex problems of chemical sensing with similar success as animals routinely do.

This special issue has the ambitious goal of bringing together biologists and engineers to report on biological solutions and engineering approaches to chemical sensing challenges in order to better understand in what aspects both fields can find common ground of discussion and to thus promote novel areas of interdisciplinary research.

Table of Contents

05 ***Bio-Inspired Solutions to the Challenges of Chemical Sensing***
Ramon Huerta and Thomas Nowotny

07 ***Sequential Mechanisms Underlying Concentration Invariance in Biological Olfaction***
Thomas A. Cleland, Szu-Yu T. Chen, Katarzyna W. Hozer, Hope N. Ukatu, Kevin J. Wong and Fangfei Zheng

19 ***Glomerular Latency Coding in Artificial Olfaction***
Jaber Al Yamani, Farid Boussaid, Amine Bermak and Dominique Martinez

28 ***An Investigation on the Role of Spike Latency in an Artificial Olfactory System***
Eugenio Martinelli, Davide Polese, Francesca Dini, Roberto Paolesse, Daniel Filippini, Ingemar Lundström and Corrado Di Natale

40 ***Synaptic Inhibition Controls Transient Oscillatory Synchronization in a Model of the Insect Olfactory System***
Collins Assisi and Maxim Bazhenov

50 ***Detailed Characterization of Local Field Potential Oscillations and Their Relationship to Spike Timing in the Antennal Lobe of the Moth Manduca Sexta***
Kevin C. Daly, Roberto F. Galán, Oakland J. Peters and Erich M. Staudacher

72 ***Parallel Representation of Stimulus Identity and Intensity in a Dual Pathway Model Inspired by the Olfactory System of the Honeybee***
Michael Schmuker, Nobuhiro Yamagata, Martin Paul Nawrot and Randolf Menzel

85 ***Olfactory Sensor Processing in Neural Networks: Lessons from Modeling the Fruit Fly Antennal Lobe***
J. Henning Proske, Marco Wittmann and C. Giovanni Galizia

97 ***Non-Linear Blend Coding in the Moth Antennal Lobe Emerges from Random Glomerular Networks***
Alberto Capurro, Fabiano Baroni, Shannon B. Olsson, Linda S. Kuebler, Salah Karout, Bill S. Hansson and Timothy C. Pearce

113 ***Network Architecture Underlying Maximal Separation of Neuronal Representations***
Ron A. Jortner

138 ***Characterization of a Clinical Olfactory Test with an Artificial Nose***
David J. Yáñez, Adolfo Toledano, Eduardo Serrano, Ana M. Martín de Rosales, Francisco B. Rodríguez and Pablo Varona

143 ***Sensor Selection and Chemo-Sensory Optimization: Toward an Adaptable Chemo-Sensory System***
Alexander Vergara and Eduard Llobet

164 Mobile Robots for Localizing Gas Emission Sources on Landfill Sites: Is Bio-Inspiration the Way to Go?
Victor Hernandez Bennetts, Achim J. Lilienthal, Patrick P. Neumann and Marco Trincavelli

176 Evolving a Neural Olfactorimotor System in Virtual and Real Olfactory Environments
Paul A. Rhodes and Todd O. Anderson

frontiers in
NEUROENGINEERING

EDITORIAL
published: 29 October 2012
doi: 10.3389/fneng.2012.00024

Bio-inspired solutions to the challenges of chemical sensing

Ramon Huerta*[1] *and Thomas Nowotny*[2]*

[1] *BioCircuits Institute, University of California, San Diego, CA, USA*
[2] *CCNR, Informatics, University of Sussex, Brighton, UK*
**Correspondence: t.nowotny@sussex.ac.uk*

Edited by:
Laura Ballerini, University of Trieste, Italy

Reviewed by:
Laura Ballerini, University of Trieste, Italy

It is widely recognized that further breakthroughs in science and technology may rely on multidisciplinary research efforts. Breaking the boundaries of well-established research fields and combining methodologies from disparate areas can foster innovative and translational research. Chemical sensing is no exception. Perhaps more than any other sensory modality, chemical sensing is plagued with major technical and conceptual challenges: the turbulent nature of the signal carrier, the long term instability, lack of sensitivity, and slow response times of sensors and the lack of a reliable odor map to characterize mammalian perception. When facing these hurdles, the designers of artificial devices for gas recognition look at the olfactory system of animals for inspiration because animals seemingly effortlessly accomplish some of the unsolved challenging problems in machine olfaction: recognizing odors and odor mixtures from a chemical background, segmenting mixtures of odors into components, being sensitive and robust and extracting the same odor percept over a wide range of concentrations.

Our challenge in bio-mimetic chemical sensing is to identify at all levels from the olfactory receptors to the central nervous system, what are the key ingredients to these impressive abilities. The goal of this research topic is to document highlights from this ongoing effort and to compile an up-to-date overview not only from the academic point of view but also with respect to industrial applications.

This research topic emanates from our own effort to bridge the anatomical and physiological data of the olfactory system, in particular in insects, to explain pattern recognition (Huerta et al., 2004; Nowotny et al., 2005; Huerta and Nowotny, 2009) and apply them to real problems with artificial sensor arrays and other applications (Muezzinoglu et al., 2008, 2009; Huerta et al., 2012). This research topic therefore brings together researchers from chemistry, neuroscience, physics, biology, and computer science, and the described work extends from fundamental scientific questions to technological applications.

On the scientific side the contributions tackle three core issues: concentration-invariant representations of odors, properties, and the potential role of oscillations in the olfactory system and the nature of odor interactions in mixtures.

Cleland et al. (2012) address concentration-invariant odor perception in rats and find that there are six known mechanisms that combine to achieve odor representations that do only minimally depend on concentration. Yamani et al. (2012) take a different view on concentration-invariant odor perception. Taking inspiration from the convergence of olfactory receptor neurons onto glomeruli and the use of latency as the coding signal they design a bio-mimetic information processing method for a metal oxide gas sensor array. Martinelli et al. (2011) on the other hand have identified another advantage of latency coding. They propose a bio-inspired solution to accelerate the response to odors using a network that uses the spike latency to discriminate volatile gases.

The presence of clear but transient oscillations in many olfactory systems has baffled scientists for a long time. The review of Assisi and Bazhenov (2012) summarizes their extensive work on the origin of these oscillations. However, Daly et al. (2011) report that they discovered that in the hawk moth *Manduca Sexta* oscillatory patterns have quite different characteristics than previously reported: oscillations are frequency modulated by odor input rather than amplitude-modulated, they are local rather than global and spikes lock to the oscillations during baseline activity more than during odor stimuli. Schmuker et al. (2011) address the problem of concurrent odor recognition and odor concentration estimation using models of the bee olfactory system and demonstrate that either goal can be achieved by the same antennal lobe network depending on the strength of lateral inhibition. In a similar direction, Proske et al. (2012) investigate different network topologies in an antennal lobe model of the fruit fly *Drosophila* and find that networks with a heterogeneous, correlation based inhibitory network lead to the best discrimination performance. Capurro et al. (2012) address the problem of odor interactions in odorant blends. Comparing experimental data from local and projections neurons with a reduced model they conclude that the experimentally observed odor interactions can be explained solely through local neuron and projection neuron response profiles and a simple mechanism of lateral inhibition.

On the technological side we are pleased to have received a variety of articles describing applications of chemical sensors and electronic noses, robotics, and a clinical application for hyposmia/anosmia characterization (Yánez et al., 2012).

Vergara and Llobet (2012) address the problem of better selectivity on sensor technology using signal modulation which is directly related to the oscillatory behavior investigated in the

biological systems above. In robotic implementations of gas sensing, Hernandez Bennetts et al. (2012) compare biologically plausible models and statistical strategies for odor localization. Their results are not too optimistic because, as they argue, current sensor technology still differ fundamentally in their sensing and actuation capabilities from highly sensitive and fast biological chemical receptors. It is clear that we still need to develop more sensitive, faster, and more stable sensors to be able to take full advantage of bio-inspired technology here. Nevertheless, in the meantime, until this technology is fully developed there are alternative hybrid approaches that allow testing strategies with simulated gas sensor responses and real robots as described in this research topic by Rhodes and Anderson (2012). This idea may point out a direction in which multidisciplinary approaches can advance in seeking the prospects of the neuro-inspired algorithms for information processing.

In summary this special issue is providing an overview of the frontiers in olfactory processing in the brain and the current challenges in applying these neuro-inspired principles in artificial olfaction. Much progress has been made, and it is becoming clear that the path to future breakthroughs does not emerge from isolated fields, but from multidisciplinary efforts like the ones presented here.

ACKNOWLEDGMENTS

Ramon Huerta acknowledges support by NIDCD, R01DC011422-01. Thomas Nowotny was partially supported by an RCUK Academic Fellowship.

REFERENCES

Assisi, C., and Bazhenov, M. (2012). Synaptic inhibition controls transient oscillatory synchronization in a model of the insect olfactory system. *Front. Neuroeng.* 5:7. doi: 10.3389/fneng.2012.00007

Capurro, A., Baroni, F., Olsson, S. B., Kuebler, L. S., Karout, S., Hansson, B. S., et al. (2012). Non-linear blend coding in the moth antennal lobe emerges from random glomerular networks. *Front. Neuroeng.* 5:6. doi: 10.3389/fneng.2012.00006

Cleland, T. A., Chen, S. T., Hozer, K. W., Ukatu, H. N., Wong, K. J., and Zheng, F. (2012). Sequential mechanisms underlying concentration invariance in biological olfaction. *Front. Neuroeng.* 4:21. doi: 10.3389/fneng.2011.00021

Daly, K. C., Galán, R. F., Peters, O. J., and Staudacher, E. M. (2011). Detailed characterization of local field potential oscillations and their relationship to spike timing in the antennal lobe of the moth *Manduca sexta*. *Front. Neuroeng.* 4:12. doi: 10.3389/fneng.2011.00012

Hernandez Bennetts, V., Lilienthal, A. J., Neumann, P. P., and Trincavelli, M. (2012). Mobile robots for localizing gas emission sources on landfill sites: is bio-inspiration the way to go? *Front. Neuroeng.* 4:20. doi: 10.3389/fneng.2011.00020

Huerta, R., and Nowotny, T. (2009). Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural Comput.* 21, 2123–2151.

Huerta, R., Nowotny, T., Garcia-Sanchez, M., Abarbanel, H. D. I., and Rabinovich, M. I. (2004). Learning classification in the olfactory system of insects. *Neural Comput.* 16, 1601–1640.

Huerta, R., Vembu, S., Amigó, J. M., Nowotny, T., and Elkan, C. (2012). Inhibition in multiclass classification. *Neural Comput.* 24, 2473–2507.

Martinelli, E., Polese, D., Dini, F., Paolesse, R., Filippini, D., Lundström, I., et al. (2011). An investigation on the role of spike latency in an artificial olfactory system. *Front. Neuroeng.* 4:16. doi: 10.3389/fneng.2011.00016

Muezzinoglu, M. K., Huerta, R., Abarbanel, H. D. I., Ryan, M. A., and Rabinovich, M. I. (2009). Chemosensor-driven artificial antennal lobe transient dynamics enable fast recognition and working memory. *Neural Comput.* 21, 1018–1037.

Muezzinoglu, M. K., Vergara, A., Huerta, R., Nowotny, T., Rulkov, N., Abarbanel, H. D. I., et al. (2008). Artificial olfactory brain for mixture identification. *Adv. Neural Inf. Process. Syst.* 22, 1021–1028.

Nowotny, T., Huerta, R., Rabinovich, M. I., and Abarbanel, H. D. I. (2005). Self-organization in the olfactory system: one shot odor recognition in insects. *Biol. Cybern.* 93, 436–446.

Proske, J. H., Wittmann, M., and Galizia, C. G. (2012). Olfactory sensor processing in neural networks: lessons from modeling the fruit fly antennal lobe. *Front. Neuroeng.* 5:2. doi: 10.3389/fneng.2012.00002

Rhodes, P. A., and Anderson, T. O. (2012). Evolving a neural olfactorimotor system in virtual and real olfactory environments. *Front. Neuroeng.* 5:22. doi: 10.3389/fneng.2012.00022

Schmuker, M., Yamagata, N., Nawrot, M. P., and Menzel, R. (2011). Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee. *Front. Neuroeng.* 4:17. doi: 10.3389/fneng.2011.00017

Vergara, A., and Llobet, E. (2012) Sensor selection and chemo-sensory optimization: toward an adaptable chemo-sensory system. *Front. Neuroeng.* 4:19. doi: 10.3389/fneng.2011.00019

Yamani, J. A., Boussaid, F., Bermak, A., and Martinez, D. (2012). Glomerular latency coding in artificial olfaction. *Front. Neuroeng.* 4:18. doi: 10.3389/fneng.2011.00018

Yáñez, D. J., Toledano, A., Serrano, E., Martín de Rosales, A. M., Rodríguez, F. B., and Varona, P. (2012). Characterization of a clinical olfactory test with an artificial nose. *Front. Neuroeng.* 5:1. doi: 10.3389/fneng.2012.00001

Received: 08 October 2012; accepted: 10 October 2012; published online: 29 October 2012.

*Citation: Huerta R and Nowotny T (2012) Bio-inspired solutions to the challenges of chemical sensing. Front. Neuroeng. **5**:24. doi: 10.3389/fneng.2012.00024*

frontiers in
NEUROENGINEERING

HYPOTHESIS AND THEORY ARTICLE
published: 05 January 2012
doi: 10.3389/fneng.2011.00021

Sequential mechanisms underlying concentration invariance in biological olfaction

Thomas A. Cleland*, Szu-Yu T. Chen, Katarzyna W. Hozer, Hope N. Ukatu, Kevin J. Wong and Fangfei Zheng

Computational Physiology Laboratory, Department of Psychology, Cornell University, Ithaca, NY, USA

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Maxim Bazhenov, University of California, USA
Barbara Webb, University of Edinburgh, UK

****Correspondence:***
Thomas A. Cleland, Computational Physiology Laboratory, Department of Psychology, Cornell University, Ithaca, NY 14853, USA.
e-mail: tac29@cornell.edu

Concentration invariance—the capacity to recognize a given odorant (analyte) across a range of concentrations—is an unusually difficult problem in the olfactory modality. Nevertheless, humans and other animals are able to recognize known odors across substantial concentration ranges, and this concentration invariance is a highly desirable property for artificial systems as well. Several properties of olfactory systems have been proposed to contribute to concentration invariance, but none of these alone can plausibly achieve full concentration invariance. We here propose that the mammalian olfactory system uses at least six computational mechanisms in series to reduce the concentration-dependent variance in odor representations to a level at which different concentrations of odors evoke reasonably similar representations, while preserving variance arising from differences in odor quality. We suggest that the residual variance then is treated like any other source of stimulus variance, and categorized appropriately into "odors" via perceptual learning. We further show that naïve mice respond to different concentrations of an odorant just as if they were differences in quality, suggesting that, prior to odor categorization, the learning-independent compensatory mechanisms are limited in their capacity to achieve concentration invariance.

Keywords: learning, odor representations, categorization, generalization, olfactory bulb, concentration invariance, mice, computational neuroscience

INTRODUCTION

In natural environments, odorants (analytes) can vary over many orders of magnitude in concentration—from ripe fruit or carrion in close proximity to the subtle scents of a trail of secretions or distant prey. In order to recognize known odors across the ranges of concentration at which they may be encountered, the olfactory system must in some way achieve *concentration invariance* in its odor representations, somehow separating concentration-dependent effects from information representing odor quality so that the odor source can be correctly identified.

Intensity invariance is a common problem across sensory systems, largely because the physical properties of the external environment vary to a much wider extent than the limited dynamic ranges of primary sensory receptors are able to capture. However, the problem is particularly acute in chemosensory modalities. Like all sensory receptors, primary chemosensors exhibit broad receptive fields that respond differentially to changes in intensity (concentration) as well as to changes in stimulus quality. Additionally, however, increasing odorant concentrations also recruit novel, lower-affinity ligand-receptor interactions that can interfere in unpredictable ways with existing interactions. The net effect is that, in addition to relatively predictable monotonic changes in receptor activation levels, concentration changes affect odor representations in unpredictable ways that are indistinguishable from changes in odor quality. Indeed, these effects essentially *are* changes in odor quality, as they arise from qualitative changes in the pattern of ligand-receptor interactions across the olfactory epithelium (**Figure 1**); interestingly, some odorants are perceived to shift in quality more than others when presented at different concentrations (Gross-Isseroff and Lancet, 1988; Johnson and Leon, 2000; Wright et al., 2005). The problem of olfactory concentration invariance consequently has been of considerable and persistent interest (Gross-Isseroff and Lancet, 1988; Duchamp-Viret et al., 1990; Bhagavan and Smith, 1997; Cleland and Linster, 2002; Cleland and Narla, 2003; Stopfer et al., 2003; Cleland et al., 2007; Uchida and Mainen, 2007).

Animals, including humans, are able to recognize many odors across reasonably wide ranges of concentration, and this capacity is critical for the utility of artificial noses as well. What algorithms underlie this capability of biological olfactory systems, and how can they be adapted to artificial systems? To date, several alternatives have been explored, many with considerable merit, though none plausibly achieve the nominal goal of concentration invariance in its entirety. We here argue that true concentration invariance in chemosensory systems is not achievable, and instead outline a practical if imperfect solution to the problem that is effective in biological systems and calls for specific design elements in biomimetic artificial systems. Specifically, the unpredictable effects of concentration-dependent variance at the ligand-receptor interface produce a lossy, and therefore irreversible, transformation in the representation of a given odor across concentrations that is indistinguishable from the effects of quality-dependent variance. Rather than attempt solely to unravel the respective effects of concentration and quality on odor

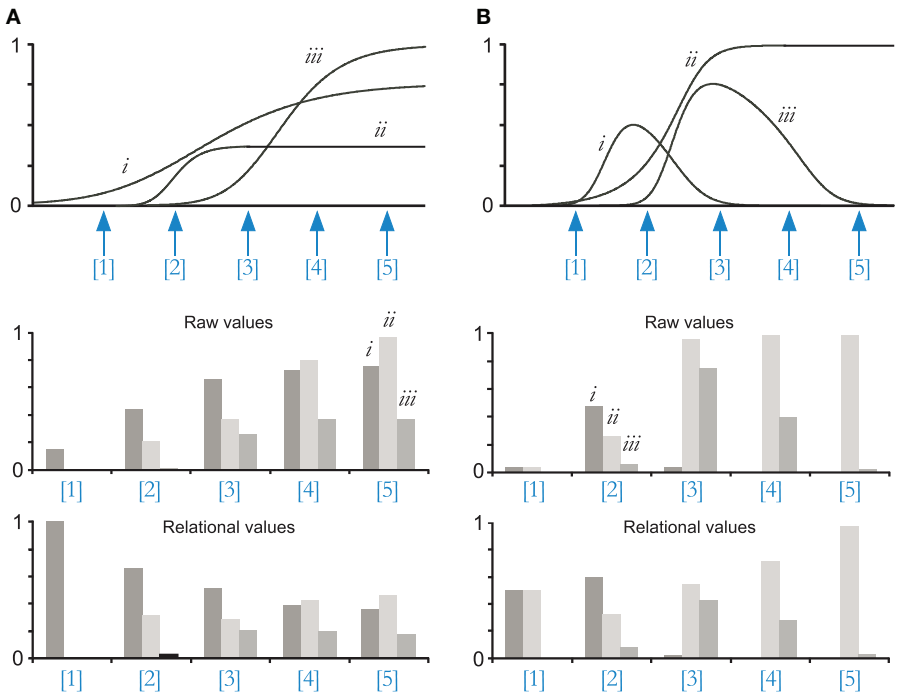

FIGURE 1 | Depiction of the problem of concentration invariance. (A) Simple models of concentration invariance are predicated upon the principle that increases in concentration generate predictably monotonic increases in the activation levels of all sensitive receptors. The broad aggregate dose-response curves of glomeruli, hypothesized to combine inputs from similarly tuned OSNs that exhibit different half-activation concentrations owing to differences in receptor reserve, can in principle extend this quasi-linear range and thereby improve the similarity of relational representations of odorants across concentrations. *Top panel*. In a computational model of ligand-receptor interactions, three ORs are activated by three odotopes of an odorant presented at a range of concentrations (five of which are labeled: [1]–[5]). Dose-response curves that do not rise to a maximum value of 1 connote odotopes that are partial agonists for their cognate ORs. Ligand-receptor interaction *i* exhibits a glomerular Hill equivalent [exponent of the population dose-response function; (Cleland and Linster, 1999)] of 0.2, yielding a quasi-linear dose-response range extending across roughly five orders of magnitude in concentration. Interactions *ii* and *iii* exhibit somewhat higher—i.e., less extreme—Hill equivalents in this example and hence have steeper, narrower dose-response curves. As a result of these broadened curves, the relational representation of the odorant across concentrations is recognizable to some degree across modest concentration ranges. *Middle panel*. Primary odor representations at five concentrations, directly read as activation levels at each of the three OR interactions depicted (identified on graph of concentration [5]). *Lower panel*. Data from the middle panel, divisively normalized so that the activity resulting from each odor presentation sums to a constant. Odor representations at concentrations [4] and [5], and to some extent [3], are reasonably similar. This similarity across concentrations will improve if the quasi-linear ranges of OR interactions *ii* and *iii* are extended to resemble that of interaction *i*. **(B)** *Top panel*. Allosteric and other non-competitive interactions, even low-affinity interactions, can render dose-response profiles at individual ORs non-monotonic, generating variance that cannot be resolved by broadening glomerular intensity tuning ranges. Adding low-affinity non-competitive interactions to the model generated clearly non-monotonic dose-response profiles for odotopes *i* and *iii*. *Middle panel*. Primary odor representations at five concentrations, directly read as activation levels at each of the three OR interactions depicted (identified on graph of concentration [2]). *Lower panel*. Data from the middle panel, divisively normalized so that the activity resulting from each odor presentation sums to a constant. Odor representations are unrecognizable across even similar concentrations, even after normalization.

representations, we propose that a series of design features in the vertebrate olfactory system serve to (a) extend the capacity of the system to represent variance within a monotonic, quasi-linear regime, (b) utilize this capacity to reduce the magnitude of variance attributable to concentration changes when possible, and finally (c) categorize the range of odorant representations attributable to concentration series together via the same learning process by which variance in odor quality is categorically grouped so as to form odors. In essence, different concentrations of a given odorant are treated as a range of reasonably similar odors that can come to be interpreted as the same odor via learning. This proposed series of design features includes, in order: (1) adaptive sampling behaviors, (2) expansion of the quasi-linear range of olfactory receptor dose-response curves via receptor reserve and axonal convergence, (3) compression of this broad intensity tuning range into the modest dynamic range of single neurons, (4) dynamical matching of pre and postsynaptic dose-response curves at the first synapse in order to optimize the detection of small changes in the chemosensory environment, (5) relational normalization, the first stage at which there is competition among different chemosensors, and finally (6) generalization of the odor representation across concentrations by categorical learning.

THEORETICAL FOUNDATIONS

PRINCIPLES OF ODOR LIGAND-RECEPTOR INTERACTION

Primary olfactory sensory neurons (OSNs) in the vertebrate nasal cavity each canonically express one or a few species of

odorant receptor proteins (ORs). In mice, there are over 1000 different functional odorant receptors expressed (Mombaerts, 2004), and millions of OSNs in the nasal cavity. The thousands of OSNs that express any given OR complement are distributed around the nasal cavity (Schoenfeld and Cleland, 2005, 2006), but their axonal projections converge and arborize together to form discrete target locations on the superficial olfactory bulb known as *glomeruli* (**Figure 2**). Setting aside the fact that glomeruli are anatomically duplicated—most glomeruli are replicated on the medial and lateral aspects of each olfactory bulb, yielding four glomeruli in total per OR per animal; (Schoenfeld and Cleland, 2005)—there is a direct correspondence between a given glomerulus, the OR complement expressed by its constituent OSNs, and the primary chemoreceptive field expressed by that OR complement (Belluscio et al., 2002; Treloar et al., 2002).

The chemoreceptive fields of individual ORs are measurable (Araneda et al., 2000), but ultimately unknowable in practice for computational purposes. A complete functional characterization of an OR would require not only an exhaustive list of strong and moderate agonists, but also a comparably exhaustive characterization of odorants acting as weak agonists, antagonists, and non-competitive allosteric modulators that would by their presence impair that OR's capacity to respond to odorant agonists. Interfering ligands in this sense need not be odorants arising from distinct sources, or even different molecules within odor mixtures (which comprise most natural odors), but also can arise from multiple receptor binding sites presented by individual odorant

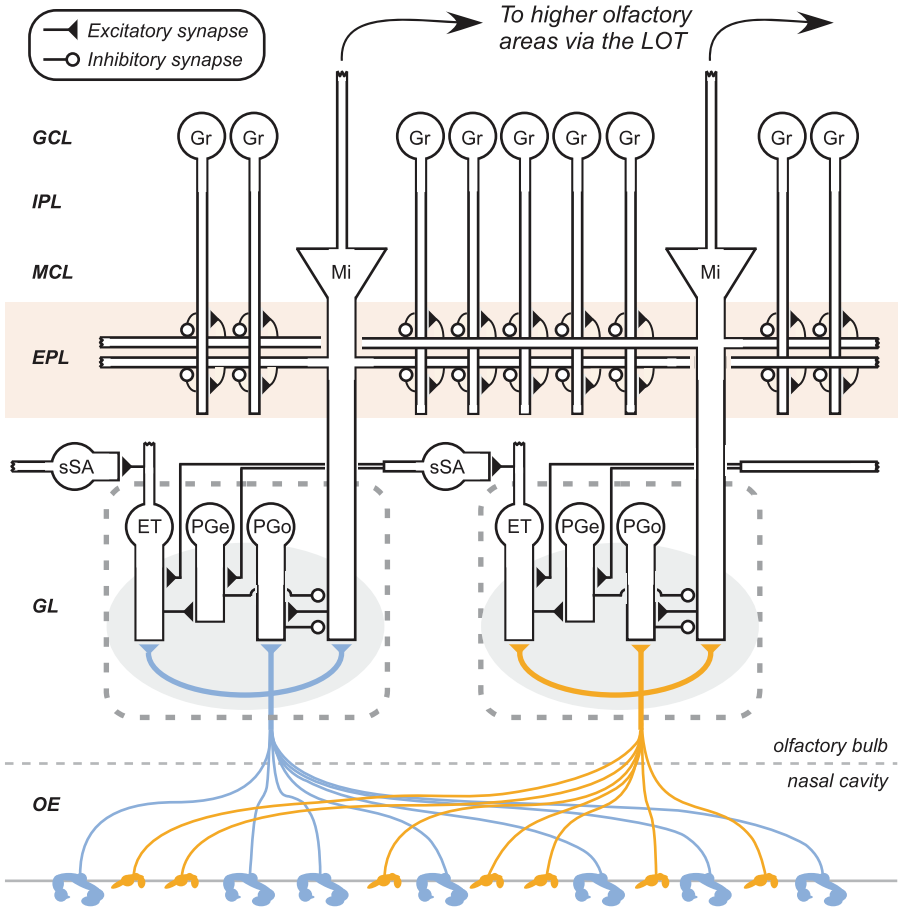

FIGURE 2 | Circuit diagram of the mammalian olfactory bulb (two glomeruli shown, with corresponding postglomerular circuitry). The axons of olfactory sensory neurons (OSNs) expressing the same odorant receptor type (denoted by the shape and color of the receptor) converge together to form glomeruli (shaded ovals) on the surface of the olfactory bulb. Multiple classes of olfactory bulb neuron also innervate each glomerulus. Glomerular interneuron classes are heterogeneous, and include olfactory nerve-driven periglomerular cells (PGo), external tufted cell-driven periglomerular cells (PGe), and multiple subtypes of external tufted cells (ET). Superficial short-axon cells (sSA) are not associated with specific glomeruli but project broadly and laterally within the deep glomerular layer, interacting with glomerular interneurons. Principal neurons include mitral cells (Mi), which interact via reciprocal connections in the external plexiform layer (EPL) with the dendrites of inhibitory granule cells (Gr), thereby receiving recurrent and lateral inhibition. Middle/deep tufted cells, another class of olfactory bulb principal neurons, are not depicted. OE, olfactory epithelium (in the nasal cavity); GL, glomerular layer; EPL, external plexiform layer; MCL, mitral cell layer; IPL, internal plexiform layer; GCL, granule cell layer. Filled triangles denote excitatory (glutamatergic) synapses; open circles denote inhibitory (GABAergic) synapses. Speckles surrounding OSN terminals connote volume-released GABA and dopamine approaching presynaptic $GABA_B$ and dopamine D2 receptors. Figure adapted from (Cleland, 2010).

molecules. In the relatively unregulated chemical environment to which the nasal cavity is subjected, these complex pharmacological interactions can substantially impact the net activation levels of OSNs, and therefore alter the primary odor representation.

ODOR SPACE

As the competitive and non-competitive ligand-receptor interactions underlying primary olfactory transduction are unknowably complex, models of olfactory similarity space—commonly referred to as "odor space" and analogous to the one-dimensional space of pitch similarity in audition or the two-dimensional space of retinotopic location in early vision—cannot be usefully based on the physicochemical properties of odorants. Whereas in audition the tonotopic distribution of frequency sensitivities across the lengths of the cochlea or cochlear nucleus subdivisions can be mapped with respect to physical frequency itself (Luo et al., 2009), the analogous relationship in olfaction between OSN activation levels and the physicochemical properties of odorant molecules cannot practically be systematically measured. However, the pattern of OSN activation levels itself is a perfectly valid, albeit species-specific, foundational metric on which to base analyses of odor similarity. In this metric, any given olfactory representation (at a point in time) is uniquely defined as the pattern of levels of activation across each of the ∼1000 (in mice) ORs, with no reference made to the many possible configurations of agonists, antagonists, and allosteric site ligands that could underlie that pattern. (Indeed, there is by definition no way for the olfactory system to distinguish among different ligand configurations that result in the same pattern of input activity, so from the perspective of olfactory coding no information is lost). Moreover, the elemental odor stimulus is defined here not as an odorant molecule or molecular epitope *per se*, but rather an *odotope*, here defined as "the net effect of a given odorant molecule on a single type of odorant receptor" without direct correspondence to molecular structural features (Cleland, 2008). Any changes in the pattern of chemoreceptor activation consequently are treated equally, whether deriving from differences in concentration or in ligand complement.

In this framework, odor space is concretely defined as having a dimensionality equal to the number of different ORs—in mice, roughly 1000—because each OR can in principle be activated independently of any other. Conveniently, these dimensions directly correspond to olfactory bulb glomeruli and hence can be directly observed using various experimental techniques (Friedrich and Korsching, 1997; Belluscio and Katz, 2001; Leon and Johnson, 2003; Fletcher et al., 2009). Theoretically, any given odor representation can at a given point in time be described by an n-dimensional vector, where n is the number of different ORs in the system, and the magnitude in each dimension ranges from 0 (no activity) to 1 (maximal activation). However, the utility of this vector model is limited by the problem of variance.

VARIANCE

Variance in stimulus quality is inescapable. Even subsequent presentations of the same odorant under experimentally controlled conditions will not evoke exactly the same pattern of neural activity. Rather, each of these different evoked activity patterns constitutes a different n-dimensional vector; however, these vectors can be bound together into a common perceptual quality by virtue of their overall *similarity*. That is, odor representations are not single vectors, but n-dimensional *clouds* of vectors in odor space with characteristic sizes and shapes. The probabilistic boundaries that define the size and shape of such clouds define the region within which a meaningful odor (such as "apple") will be recognized irrespective of within-category variability (e.g., cultivar, ripeness, temperature, growing-season variables). Outside of these boundaries, an odor stimulus will be judged as to some degree different from that representation in its quality or implications. Quantitatively, these clouds constitute n-dimensional probability density functions (PDFs) that correspond to what might be termed *odors*, as distinct from odorants: learned ranges of chemosensory activation patterns that convey the same meaning. Importantly, the degree of tolerable variance in each dimension is an integral part of the odor representation; large changes in some dimensions may be included in the same odor representation whereas small differences in other dimensions may indicate a different odor with different implications.

These odor PDFs can be behaviorally measured using *generalization gradients* [**Figure 3**; (Shepard, 1987; Linster and Hasselmo, 1999; Tenenbaum and Griffiths, 2001; Cleland et al., 2002)]. Olfactory generalization gradients are systematically regulated by learning, and directly measure how progressively increasing dissimilarity among odors yields a corresponding decline in animals' expectation of similar outcomes (Daly et al., 2001; Wilson and Stevenson, 2006; Cleland et al., 2009; Fernandez et al., 2009). Olfactory generalization gradients in vertebrates also are regulated by extrinsic neuromodulation within olfactory bulb and piriform cortex (Linster and Cleland, 2002; Wilson et al., 2004; Mandairon et al., 2006; Mandairon and Linster, 2009); the level of activity in ascending neuromodulatory inputs reflects behavioral state and may underlie animals' capacity to alter the perception of similarity in accordance with task demands.

LEARNING ABOUT ODOR VARIANCE

Multiple types of olfactory generalization gradient can be used to measure the sizes and shapes of odor representations (Cleland et al., 2002). In *non-associative* generalization gradients [also known as cross-habituation or spontaneous discrimination gradients; (Cleland et al., 2002; Mandairon et al., 2006)], an animal is repeatedly presented with an odorant until its investigation time drops to an asymptotic minimum. Presentation of highly dissimilar odorants will still evoke a full investigation response, whereas odorants similar to the habituated odorant evoke partial responses. The function of investigation time with respect to odorant similarity defines the probabilistic boundary of the odor representation. *Associative* generalization gradients are similar in principle, but are measured by conditioning an animal to work (dig) for a reward when an odor cue is delivered. Presentation of odorants dissimilar to the conditioning odorant evokes no conditioned response, whereas presentation of perceptually similar odorants elicits a digging response that declines in perseverance as the odorant cue becomes more

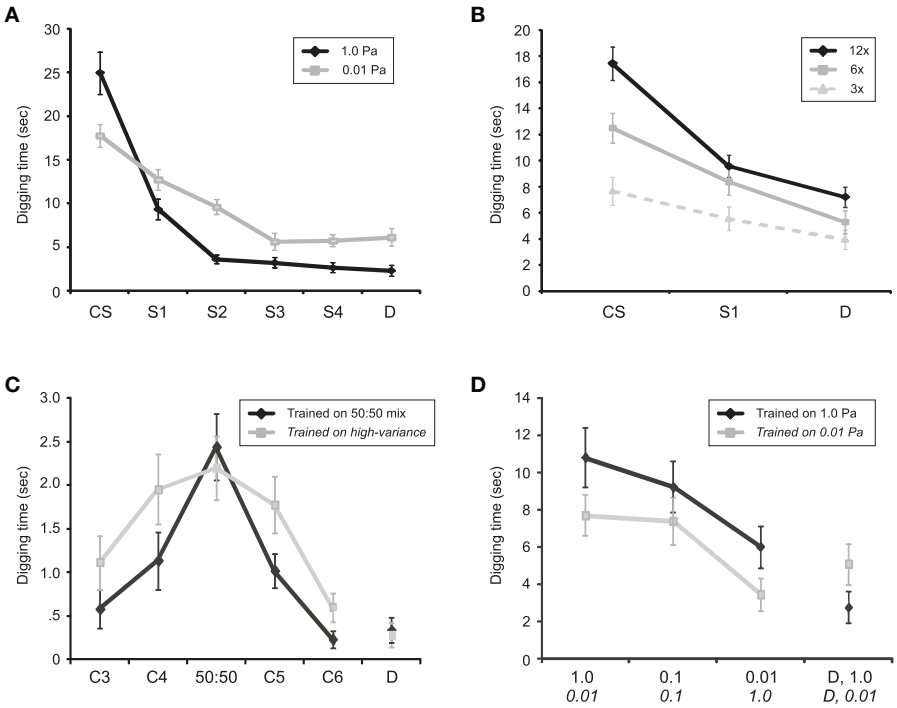

FIGURE 3 | Olfactory generalization gradients in mice. (A) Associative generalization from a conditioned odorant stimulus (CS) to a series of four sequentially similar odorants (S1-S4) plus one structurally and perceptually dissimilar control odorant (D). Presenting all odorants at a higher concentration (theoretical vapor-phase partial pressure of 1.0 Pa; *black line*) yielded a steeper, narrower generalization gradient than did identical training with low-concentration odorants (0.01 Pa; *gray line*), reflecting the learning-theoretic principle that higher CS salience supports greater learning. Twelve training trials were administered prior to testing. Figure adapted from Cleland et al. (2009). **(B)** Increasing the number of training trials (CS-reward pairings) prior to testing progressively increased perseverance and sharpened associative generalization gradients. 3×: three training trials; 6×: six training trials; 12×: 12 training trials. Figure adapted from Cleland et al. (2009). **(C)** Generalization gradients adapt to the variance of the conditioning odor. The high-variance conditioning group (see Methods) generalized fully across the range of CS variability (no difference in digging times between 50:50 and either C4 or C5; Welch test, $t(46.36) = 0.444$, $p = 0.659$; $t(47.43) = 0.854$, $p = 0.398$, respectively), whereas the low-variance group clearly distinguished both C4 and C5 from the 50:50 odor mixture CS (significant differences in digging times; Welch test, $t(56.87) = 2.583$, $p = 0.012$; $t(43.45) = 3.314$, $p = 0.002$, respectively). **(D)** Mice perceive sufficiently different concentrations of novel odorants as distinct odors. One group of mice was conditioned to an odorant CS at a high concentration (1.0 Pa; *black line*, test concentrations in Pa listed in Roman font on x-axis) and tested on two lower concentrations of that odorant as well as a dissimilar control odorant (D) at 1.0 Pa. A second group was conditioned to the same odorant CS at a low concentration (0.01 Pa; *gray line*, test concentrations in Pa listed in italic font on x-axis) and tested on two higher concentrations of that odorant as well as a dissimilar control odorant (D) at 0.01 Pa. Both groups treated the test odorant that was two orders of magnitude higher or lower in concentration as a distinct odor, roughly comparable in similarity to a structurally dissimilar odorant D. See the *Learning-Dependent Construction of Odor Representations* section for analysis details. Odor sets and vol/vol dilutions are detailed in **Table 1**. In all figures, error bars denote the standard error of the mean.

dissimilar from the conditioning odorant [**Figures 3A,B**; (Linster and Hasselmo, 1999; Cleland et al., 2002, 2009)]. The function of perseverance (e.g., digging time) with respect to odorant similarity yields a generalization gradient that delineates the *consequential region* (Shepard, 1987) of the underlying odor representation.

Critically, the shape of olfactory generalization gradients is modified by learning. Classical learning determinants such as conditioned stimulus salience and unconditioned stimulus valence systematically modify associative generalization gradients, as do changes in the numbers of training trials prior to testing (Cleland and Narla, 2003; Cleland et al., 2009). That is: increased learning in these studies corresponds to progressively sharper generalization gradients (**Figures 3A,B**; adapted from Cleland et al., 2009). Considering these generalization gradients as PDFs of odor quality—that is, maps of gradually declining probability that increasingly dissimilar odor stimuli have the same implications as the conditioned odorant—this learning effect directly reflects the statistical principle that the standard error of the mean is reduced as the number of samples increases. In a broader sense, learned generalization gradients reflect the fundamental principle of sensory representation described above: generalization gradients are reflections of perceptual learning, in which the olfactory system progressively adapts its internal representation of odor space, through experience, to the statistics of odor encounters and consequences in the world (Wilson and Stevenson, 2006; Fernandez et al., 2009; Wright et al., 2009).

A corollary of this principle is that the different primary odor representations arising from the range of odor qualities that

correspond to a meaningful common odor such as "apple" should be bound together into the same odor PDF so as to comprise a common percept. In a generalization framework, this connotes a region of full generalization across the range of odor qualities that exhibits no significant dissimilarity-dependent reduction in the operant response. This principle of generalization across similar stimulus qualities in a continuous metric space has been theoretically described in a Bayesian framework (Tenenbaum and Griffiths, 2001). We tested the predictions of this theory in a study in which two groups of mice were trained to the same degree with the same mean conditioned odorant, but with different degrees of variance in odor quality across training trials. As predicted, we observed that the full-generalization window of the olfactory generalization gradient expanded to contain all of the conditioned variance, and that generalization began to fall off just outside of the boundaries of stimulus quality defined by the training experience (**Figure 3C**).

ADDRESSING THE PROBLEM OF CONCENTRATION INVARIANCE

Ideally, changes in odorant concentration simply would increase the activity of all sensitive receptors proportionally, such that subsequent processing could easily extract the portion of stimulus variance attributable to odor quality differences. This, of course, is not the case. The sigmoidal dose-response relationships of ligand-receptor interactions have a limited quasi-linear range, but also long asymptotic tails at the extremes such that concentration changes alter the relational pattern of receptor activation on which odor quality representations are based (**Figure 1A**). Given that the quasi-linear range of ligand-receptor interactions is strictly limited by mass action law (Cleland and Linster, 1999), these distorting effects on relational representations are substantial. Moreover, the Hill coefficients of individual odorant receptors are often substantially greater than unity (Firestein et al., 1993), exacerbating this problem; in such neurons, the range between a concentration evoking 10% activation and that evoking 90% activation (EC_{10-90}) can be substantially less than two orders of magnitude. In intact olfactory epithelia, the median EC_{05-95} dynamic range has been estimated at 0.3–1.2 log units (Rospars et al., 2000). Finally, the emergence of lower-affinity ligands interacting with activated receptors can alter dose-response curves in the same way as would a novel odor component, altering effective ligand potencies and even rendering dose-response functions non-monotonic if allosteric or other non-competitive interactions arise [**Figure 1B**; discussed further in (Cleland and Linster, 1999; Cleland, 2008)]. The inevitable consequence of these physical principles is that, while a substantial fraction of the variance in odor representations attributable to odor concentration can be minimized, changing odor concentrations nevertheless results in significant residual representational differences that are indistinguishable from changes in odorant quality. This predicts that naïve animals will respond to different odor concentrations as if they were different odorants, which indeed is the case in both mice (**Figure 3D**) and honeybees (Choudhary, 2009).

The second problem of concentration is common among sensory systems: the fact that the narrow dynamic range of ligand-receptor interactions, as constrained by mass action law, is far more limited than the range of external environmental properties that are to be measured. Indeed, to be able to resolve large differences in concentration at all is a substantial feat of biological engineering, as no individual sensor is capable of resolving such a broad concentration range. Moreover, if multiple different sensors (olfactory receptor types) with differing affinities for a given ligand were employed systematically to resolve such a broad concentration range, their otherwise unrelated chemoreceptive fields would render odor quality and concentration irrevocably intertwined, such that little or none of the variance attributable to concentration could be identified as such. While this remains a problem, it is made considerably more tractable in olfaction by the early processing strategies outlined in the *Broadening of the Aggregate Dose-Response Curves of Glomeruli*, *Intensity Compression at the First Synapse*, and *Adaptation to Background Odor Intensity* sections below.

In sum, as briefly outlined above, the solution of the biological olfactory system to the problem of concentration is likely to involve multiple coordinated processing mechanisms, which we propose include the following: (1) constraining stimulus variance somewhat by behaviorally regulating the concentrations presented to the olfactory epithelium, (2) substantially extending the quasi-linear dynamic range of each individual olfactory receptor class via receptor reserve and axonal convergence, thereby reducing the distorting effects of concentration on relational representations, and (3) compressing the absolute range of intensities exhibited by OSN populations within each glomerulus to match the dynamic ranges of postsynaptic second-order neurons. This intensity compression stage collapses the expanded quasi-linear dynamic range generated at the glomerulus (*Broadening of the Aggregate Dose-Response Curves of Glomeruli*, below), thereby eliminating a substantial fraction of the variance attributable to concentration (*Intensity Compression at the First Synapse*, below). Processing at this stage also appears to include (4) an adaptive component that adjusts to stable odor backgrounds so as to emphasize changes in the chemosensory environment rather than always being dominated by the strongest stimuli, and (5) an additional, competitive stage of normalization across the input (glomerular) layer of olfactory bulb. The remaining concentration-dependent variance, while substantially reduced in magnitude, cannot be differentiated from quality-dependent variance and, we here propose, (6) is perceived as quality variance until and unless the animal learns to categorize this variance into a single odorant.

ADAPTIVE SAMPLING BEHAVIORS

Animals are capable of modulating the intensity of sniffing behavior, potentially increasing the concentration of weak odorants and limiting the intensity of strong odors in the nasal cavity by regulating the depth and frequency of inhalation (Verhagen et al., 2007, but see Teghtsoonian et al., 1978), and possibly even manipulating odor quality to a limited extent by altering the deposition pattern of odorants onto regions of the nasal mucosa that are differentially enriched with particular ORs (Schoenfeld and Cleland, 2006). Control over sniffing may be particularly important for limiting the access of highly intense odors to the nose, as the capacity for differentiation among odors is sharply reduced when

odorants are extremely strong. The regulation of sampling behaviors (sniffing, antennal flicking) and their role in perception is an ongoing field of research (Koehl, 2006; Carey et al., 2009; Shusterman et al., 2011; Wesson et al., 2011).

BROADENING OF THE AGGREGATE DOSE-RESPONSE CURVES OF GLOMERULI

Broadening the dose-response curves of olfactory ligand-receptor interactions until the quasi-linear portions of their concentration tuning ranges extend across multiple orders of magnitude would enable a substantial fraction of concentration-dependent variance in odor representations to be quasi-linearized, such that it subsequently could be selectively removed from the representation by some form of normalization. However, the breadth of the quasi-linear region [EC_{10-90}; (Cleland and Linster, 1999)] of individual ligand-receptor binding curves is limited by mass action law to less than two orders of magnitude; moreover, the reported Hill coefficients of isolated olfactory receptor neurons (Firestein et al., 1993) suggest that many of these dose-response relationships are even narrower than they need be (based on a minimum Hill coefficient of unity).

The olfactory system appears to have utilized convergence to engineer a solution to this conundrum. Thousands of individual OSNs that express the same odorant receptor complement project their axons convergently onto single locations on the surface of the olfactory bulb to form glomeruli [indeed, a substantial fraction of the glomerular volume is made up of their axonal arborizations; (Kosaka et al., 2001; Nawroth et al., 2007)]. These convergent OSNs share the same receptive field for quality, but may be differently attuned for ranges of concentration via differences in receptor reserve (Zhu, 1993) such that the net dose-response curve *of the convergent population* can be greatly extended so as to span several orders of magnitude (Cleland and Linster, 1999). Indeed, aggregate glomerular dose-response curves, as assessed by imaging studies, can span several orders of magnitude (Friedrich and Korsching, 1997; Wachowiak et al., 2002).

This substantial broadening of glomerular dose-response relationships prior to the first synapse, when followed by a normalization process, can greatly reduce the magnitude of the variance that is attributable to concentration (**Figure 1A**). Theoretical calculations of dissimilarities among representations of the same odorant across even small concentration differences suggest that concentration-invariant perception would not be remotely possible if the broadening of the quasi-linear range of these glomerular activation functions did not enable the extraction of much of this concentration-dependent variance component (Cleland et al., 2007). Nevertheless, this process is imperfect in its linearity and also contributes nothing towards resolution of the problem of allosteric and noncompetitive interactions (**Figure 1B**).

INTENSITY COMPRESSION AT THE FIRST SYNAPSE

The broadened aggregate dose-response relationships observed in OB glomeruli enable olfactory stimuli to be encoded over broad ranges of concentration. However, as this broad dynamic range greatly exceeds the dynamic range of any one ligand-receptor relationship, it must be compressed if the full dose-response relationship is to be encodable by individual postsynaptic neurons such as mitral, tufted, and periglomerular cells (or, potentially, across small groups of such cells). Indeed, mitral and tufted cells are able to differentially respond to concentration changes across reasonably wide concentration ranges (Hamilton and Kauer, 1989; Wellis et al., 1989; Duchamp-Viret et al., 1990; Nagayama et al., 2004).

Compression of the absolute range of input intensity into a manageable dynamic range appears to be mediated in large part by feedback inhibition onto the presynaptic terminals of OSNs, mediated by presynaptic $GABA_B$ receptors that reduce calcium influx into terminals and thereby reduce transmitter release (Nawroth et al., 2007; Pirez and Wachowiak, 2008). OSN activity monosynaptically excites a subclass of GABAergic periglomerular interneurons [PGo cells; **Figure 2**; (Shao et al., 2009)] and indirectly excites another subclass of GABAergic periglomerular cells (PGe cells), both of which release GABA in the vicinity of OSN presynaptic terminals (as well as onto the dendrites of mitral cells). The resulting presynaptic inhibition substantially constrains the output of convergent OSN populations. Notably, the magnitude (weight) of this presynaptic inhibition is consistent across both weak and strong levels of afferent excitation (Pirez and Wachowiak, 2008), and is not observably activity-dependent, at least on the timescales studied to date. In principle, stable presynaptic feedback inhibition is precisely the computation required to compress a monotonic dose-response relationship into a narrower dynamic range, appropriate for the graded excitation of postsynaptic neurons. An important outstanding question is the degree to which the level (as opposed to the weight) of this presynaptic feedback inhibition is determined independently for each glomerulus versus how much it may be regulated by an average level of activation computed across many glomeruli. PGo-mediated inhibition would clearly generate the former, whereas PGe-mediated inhibition, evoked in large part via excitation by external tufted (ET) cells, may to some extent mediate the latter (Cleland et al., 2007). This latter possibility should not be confused with presynaptically mediated center-surround inhibition among glomeruli, which has been clearly ruled out (Pirez and Wachowiak, 2008).

ADAPTATION TO BACKGROUND ODOR INTENSITY

In addition to $GABA_B$ receptors, OSN terminals presynaptically express D2-type dopamine receptors (Nickell et al., 1991), which respond to transmitter released by a dopaminergic subset of periglomerular neurons (Halasz et al., 1981; Gall et al., 1987; Toida et al., 2000). Like the $GABA_B$ receptors with which they are co-expressed, presynaptic D2 receptors are inhibitory (Hsia et al., 1999; Berkowicz and Trombley, 2000; Ennis et al., 2001; Davila et al., 2003), reducing the release of glutamate from OSN terminals in glomeruli. An important functional difference between these $GABA_B$-ergic and dopaminergic feedback loops, however, is that the weight of dopaminergic feedback appears to be regulated by a running average level of afferent activation over a timescale of many minutes to hours. Specifically, sharply reducing odor stimulus levels through naris occlusion reduces dopamine levels in olfactory bulb (Brunjes et al., 1985), but also,

on a somewhat longer timescale, reduces the expression of tyrosine hydroxylase in the dopaminergic periglomerular cells of the rat and mouse olfactory bulbs (Baker, 1990; Stone et al., 1990; Baker et al., 1993; Cho et al., 1996). As tyrosine hydroxylase is the rate-limiting enzyme for dopamine synthesis, reduced expression levels reflect reductions in the weight of dopaminergic feedback inhibition, suggesting a feedback-regulated system that gravitates toward a stable homeostatic level of activity. By this model, if input to a given glomerulus is consistently weak over many minutes or hours, dopaminergic feedback inhibition will be weakened until some homeostatic target level of average afferent activity is detected among postsynaptic neurons. Similarly, if high levels of presynaptic activity persist over a similar timescale, dopaminergic feedback inhibition will gradually increase in strength until the average level of postsynaptic activity is reduced to this same homeostatic baseline. The utility of such a mechanism is that the system adapts to stable background odor levels, effectively rezeroing with respect to any background such that subtle changes in odor can be detected even within a strongly odorous environment. In essence, this imposes a timescale on olfactory perception that privileges changing signals over static background, analogous to a Pacinian corpuscle or the perception of visual motion. Additionally, this adaptive mechanism is likely to help keep the absolute range of total OSN output levels within a glomerulus adjusted to the absolute dynamic range for input to postsynaptic neurons, assuming that the homeostatic set point for OSN output is appropriate to this purpose.

This model predicts that manipulating dopamine D2 receptor activation levels in the olfactory bulb would affect the perceived concentration of odorants, which is indeed the case (Wei et al., 2006; Escanilla et al., 2009). Briefly, administration of D2 receptor agonists, either systemically or via direct infusion into olfactory bulb, reduced rats' performance levels in a concentration-sensitive odor discrimination task, whereas D2 antagonists improved their performance levels, both in a dose-dependent manner. Blocking D2 receptors also mimics the effects of olfactory deprivation on the activity of mitral/tufted cells (Wilson and Sullivan, 1995).

RELATIONAL NORMALIZATION

Normalization of outputs from all glomeruli with respect to the total activity level across the olfactory bulb in principle can preserve the relative levels of activity among activated glomeruli while keeping the total bulbar activation level roughly constant, generating the first competitive interaction among different glomerular columns and thereby forming a *relational* odor representation (Cleland et al., 2007). Competitive normalization, unlike intensity compression as described above, can generate non-monotonic dose-response functions in second-order neurons (e.g., mitral cells) as increasingly active glomerular columns outcompete less strongly activated columns, such that mitral cells innervating the outcompeted columns become less strongly activated, or even inhibited, as concentrations rise (Cleland and Sethupathy, 2006; Cleland, 2010). Indeed, increasing odorant concentrations do not generally increase spike rates in mitral cells monotonically; rather, mitral cell responses vary across concentrations in more complex ways, sometimes transitioning from excitation to inhibition, or vice versa, as odorant concentrations rise (Wellis et al., 1989; Chalansonnet and Chaput, 1998). Importantly, net activity among olfactory bulb neurons changes considerably less across concentrations than does that in OSNs, and raw maps of (predominantly presynaptic) glomerular activation in response to odorant presentation predict the perceptual similarity of odorants far less well than do the same maps after global normalization (Cleland et al., 2007). Moreover, global normalization is an important requirement for some models of olfactory contrast enhancement (Cleland and Sethupathy, 2006).

Evidence for a cellular or network mechanism that can mediate global normalization in the olfactory bulb is incomplete. A measure of constitutive GABA release from granule cells may help dampen mitral cell responses, although granule cells' inhibitory synapses onto mitral cells are electrotonically distant from the primary dendrite in which mitral cell spikes can be initiated, and there is no evidence that lateral inhibition mediated by granule cells would be sufficiently broad and unbiased to globally normalize afferent activation patterns. Indeed, measurements of intercolumnar inhibitory efficacies in olfactory bulb suggest a sparse and highly specific lateral inhibitory map (Fantana et al., 2008). Alternatively, theoretical modeling of a lateral excitatory network in the deep glomerular layer (**Figure 2**; ET and sSA cells) has illustrated a mechanism by which global normalization of bulbar activity could be effected postsynaptically, utilizing an anatomically center-surround connectivity matrix (Aungst et al., 2003) to generate a uniform level of inhibition proportional to the total input activity across the olfactory bulb and delivered onto mitral cells (Cleland et al., 2007; Cleland, 2010). The potential interactions between this $GABA_A$-dependent mechanism and the presynaptic feedback inhibition mechanisms described above remain to be explored, both theoretically and experimentally.

LEARNING-DEPENDENT CONSTRUCTION OF ODOR REPRESENTATIONS

The work described above suggests several circuit mechanisms, working sequentially and in concert, that can substantially reduce the variance in odor representations attributable to concentration—e.g., that would be generated by different concentrations of the same odorant. However, they do not and cannot suffice to provide true concentration invariance in odor representations. We here propose that the remaining variance in odor representations generated by different concentrations of the same odorant is treated in the same way as variance deriving from changes in odor quality. That is: different concentrations of the same odorant are treated as different odors until and unless the animal learns that they have similar implications and categorizes them together into a common odor representation, as depicted for ranges of similar odorants in **Figure 3C**.

This model makes two predictions. First, it predicts that naïve animals will perceive different concentrations of the same odorant as different, and will generalize between them partially or not at all. Second, in order for these different concentrations to be bound together efficiently by learning, the prior processing stages described above presumably must have removed enough concentration-dependent variance for their representations to be

at least somewhat similar. (The validity of this second prediction may vary depending on the specific odorants in question—notably, pentanal and 2-hexanone have been reported to change in perceptual quality to humans given modest concentration changes, whereas pentanoic acid, methyl pentanoate, and pentanol do not; Johnson and Leon, 2000). We here show that different odor concentrations are treated by naïve mice as if they were different odors, and that the degree of perceptual similarity decreases with increasing concentration differences (**Figure 3D**; see Methods). Specifically, for the 1.0 Pa conditioned stimulus (*black line*), independent-measures analysis of variance (ANOVA) demonstrated a significant difference in digging time among the four trials, $F(3,116) = 7.964$, $p < 0.001$. Comparing the 1.0 Pa CS to the 0.01 Pa test odorant using Tukey's honestly significant difference (HSD) criterion demonstrated a significant decline in perseverance, $p = 0.040$. Similarly, for animals conditioned with an 0.01 Pa CS (*gray line*), ANOVA again demonstrated a significant difference in digging time among the four trials, $F(3,112) = 3.353$, $p = 0.022$. Comparing the 0.01 Pa CS to the 1.0 Pa test odorant confirmed a significant decline in perseverance (Tukey's HSD, $p = 0.035$). Thus, after conditioning to an odor CS, animals treated an odorant presented at two orders of magnitude higher or lower concentration as a distinct odor. Moreover, the magnitude of the decline in generalization across these two orders of magnitude in concentration is roughly comparable to the degree of difference between the CS and a highly dissimilar control odorant D, indicating that odorants presented at these two different concentrations can be perceived as differently as two quite different odorants.

SUMMARY AND CONCLUSIONS

Concentration invariance in the olfactory system is important in that odors from natural sources vary substantially in concentration from any given vantage point, and there is obvious benefit to being able to identify such odors irrespective of this variance. Indeed, several sequential, coordinated mechanisms in the olfactory system appear to be able to reduce the impact of concentration-based changes in odor representations substantially. However, unavoidable non-linearities in signal processing and the particular problem of interference from lower-affinity ligands render complete concentration invariance unachievable. In this behavioral paradigm, the remaining concentration-dependent effects on primary odor representations are indistinguishable from quality-dependent changes that signal the presence of different odors, and indeed modestly different concentrations of a given odorant are perceived as different odors by naïve animals. However, after this preprocessing cascade, moderately different odorant concentrations are not perceived as enormously different in quality, as presynaptic imaging studies would suggest, but rather as modestly different in quality; tenfold concentration differences remain quite perceptually similar, whereas one hundred-fold concentration differences approach asymptotic dissimilarity. Critically, odor representations that are at least moderately similar in perceptual quality (i.e., ranges of odor qualities in odor space) and that predict the same outcome can be grouped together through learning.

METHODS

BEHAVIORAL PROCEDURES

Olfactory generalization gradients were measured in mice according to established procedures (Cleland et al., 2002, 2009). Briefly, age-matched cohorts of male CD-1 mice (outbred strain; Charles River Laboratories, Wilmington, MA) were shaped (trained to dig for rewards in response to odor cues) from five to eight weeks of age and subsequently employed in experiments. Mice were maintained on a shifted 12L:12D cycle; all behavioral training was conducted during their dark cycle (9:00 a.m.–9:00 p.m.). Water was continuously available; mice were food-deprived for up to 18 hours preceding each session to motivate them to obtain sucrose rewards. Mice were fed immediately after an experimental session, and were not deprived of food on two subsequent days. All procedures were performed under the auspices of a protocol approved by the Cornell University Institutional Animal Care and Use Committee.

During each conditioning trial, two sand-filled dishes were placed in a chamber; one contained a sucrose reward and was scented with the conditioned odorant CS, whereas the other contained no reward and no odorant. Each trial began when the mouse entered the chamber, at which point it encountered the dishes and was allowed to dig in both dishes until it retrieved the reward. The mouse was then removed for a one minute intertrial interval, during which the dishes were prepared for the next trial. After the training trials were complete (12 training trials, except for **Figure 3B**), the test trials were begun. In the test trials, one dish was scented either with the CS odorant or with one of a series of similar or dissimilar test odorants, whereas the other contained no odorant, and neither dish contained any reward. The amount of time that a mouse spent digging in the scented sand (perseverance) served as the dependent variable. The duration of test trials was one minute, whereas conditioning trials ended after mice recovered the sucrose reward (up to a maximum of one minute). Intertrial intervals were one minute long, and test trials began directly after the completion of the conditioning trials. In all cases, data are aggregates of multiple separate odor sets averaged together to ensure that results are not specific to a given odorant series (**Table 1**), and testing orders were randomized and counterbalanced.

Figures 3A,B are adapted from (Cleland et al., 2009); experimental details can be found therein. All odorants included in **Figure 3B** were presented at 0.01 Pa. In the experiment comprising **Figure 3C**, two groups of mice were assembled, a low-variance conditioning group and a high-variance conditioning group. A homologous series of four odorants (C3–C6, arbitrary labels) plus a dissimilar odorant D were employed; all odorants were presented at 1.0 Pa. Mice were given 12 training trials with a conditioned odorant, which was a mixture of odorants C4 and C5. Specifically, the low-variance group was trained on a 50:50 mixture of C4 and C5 for all 12 trials, whereas the high-variance group was trained on six different C4:C5 mixture ratios centered on but not including 50:50 (specifically: 95:5, 80:20, 60:40, 40:60, 20:80, 5:95, each presented twice in a randomized and counterbalanced order). Both groups then were tested using an identical set of six test odorants: the 50:50 mixture, unmixed

odorants C4 and C5, the structurally similar odorants C3 and C6, and a dissimilar control odorant D.

In the experiment comprising **Figure 3D**, two groups of mice were conditioned to an odorant CS over 12 training trials and then tested using that same odorant at three different concentrations as well as with a structurally different odorant presented at the same vapor-phase concentration as the conditioned odorant. Critically, one group was trained on a relatively high concentration (1.0 Pa; see *Odor Sets and Dilutions* below) and tested on lower concentrations, whereas the other was trained on a relatively low concentration (0.01 Pa) and tested on higher concentrations. This enabled us to rule out an alternative interpretation, derived from work in honeybees (Pelz et al., 1997), in which the higher-concentration odorant might simply comprise a superior exemplar of the odorant, such that animals' responses to a higher-concentration odorant would be stronger than those to a lower concentration irrespective of which concentration had been explicitly conditioned. Note that the absolute level of conditioning to the lower concentration is less than that to the higher concentration, as expected (compare to **Figure 3A**), so the two generalization gradients are considered separately.

ODOR SETS AND DILUTIONS

Multiple odor sets were used to enable counterbalancing among subjects and ensure that results were not dependent on the use of specific odor sets. All mice in a cohort were tested using every odor set employed in the corresponding study. Each odor set consisted of a homologous series of 2–5 structurally similar, unbranched aliphatic odorant molecules plus one structurally dissimilar odorant used as a control; one study also utilized a binary mixture of two structurally adjacent odorants, whereas another utilized multiple concentrations of single conditioned odorants (see *Behavioral Procedures* above; **Table 1**). Vapor pressures of pure odorants were estimated with the Hass-Newton equation as implemented in ACD/Boiling Point and Vapor Pressure Calculator (version 4.5; Advanced Chemistry Development, Toronto, ON, Canada); pure odorants were

Table 1 | Odorant sets with corresponding vol/vol dilutions in mineral oil.

Odorant	Dilution for 1.0 Pa	Odorant	Dilution for 1.0 Pa
FIGURE 3C		*n*-hexanoic acid	1.49×10^{-2}
Acetic acid	0.78×10^{-4}	***n*-heptanoic acid**	$\mathbf{4.63 \times 10^{-2}}$
Propanoic acid	$\mathbf{3.31 \times 10^{-4}}$	***n*-octanoic acid**	$\mathbf{13.7 \times 10^{-2}}$
***n*-butanoic acid**	$\mathbf{12.7 \times 10^{-4}}$	*n*-nonanoic acid	36.8×10^{-2}
n-pentanoic acid	45.0×10^{-4}	neryl acetate	16.4×10^{-2}
3-heptanone	6.46×10^{-4}		
		FIGURE 3D	
Propyl acetate	0.63×10^{-4}	***n*-amyl acetate**	$\mathbf{7.23 \times 10^{-4}}$
***n*-butyl acetate**	$\mathbf{2.19 \times 10^{-4}}$	anisole	5.15×10^{-4}
***n*-amyl acetate**	$\mathbf{7.23 \times 10^{-4}}$		
n-hexyl acetate	22.7×10^{-4}	***n*-butanoic acid**	$\mathbf{12.7 \times 10^{-4}}$
Anisole	5.15×10^{-4}	3-heptanone	6.46×10^{-4}
n-pentanol	0.74×10^{-3}	**2-furyl methyl ketone**	$\mathbf{2.59 \times 10^{-3}}$
***n*-hexanol**	$\mathbf{2.55 \times 10^{-3}}$	*n*-butyl *n*-butyrate	1.65×10^{-3}
***n*-heptanol**	$\mathbf{8.38 \times 10^{-3}}$		
n-octanol	26.7×10^{-3}	***n*-butyl *n*-pentanoate**	$\mathbf{5.72 \times 10^{-3}}$
2-furyl methyl ketone	2.59×10^{-3}	Citronellal	16.6×10^{-3}
n-hexanal	2.21×10^{-4}	***n*-hexanol**	$\mathbf{2.55 \times 10^{-3}}$
***n*-heptanal**	$\mathbf{7.07 \times 10^{-4}}$	neryl acetate	1.64×10^{-3}
***n*-octanal**	$\mathbf{14.7 \times 10^{-4}}$		
n-nonanal	63.2×10^{-4}	**octanal**	$\mathbf{1.47 \times 10^{-3}}$
2,3,5-trimethylpyrazine	13.9×10^{-4}	trans-2-hexenyl acetate	1.63×10^{-3}
n-butyl propanoate	0.60×10^{-3}	***n*-hexanoic acid**	$\mathbf{14.9 \times 10^{-3}}$
***n*-butyl *n*-butyrate**	$\mathbf{1.65 \times 10^{-3}}$	*n*-heptanol	8.38×10^{-3}
***n*-butyl *n*-pentanoate**	$\mathbf{5.72 \times 10^{-3}}$		
n-butyl *n*-hexanoate	16.3×10^{-3}	**hexanal**	$\mathbf{2.21 \times 10^{-4}}$
n-butyl glycidyl ether	1.85×10^{-3}	2-hexanone	1.80×10^{-4}

*Odorants were diluted in mineral oil to concentrations theoretically emitting vapor-phase partial pressures of 1.0, 0.1, or 0.01 Pa as indicated; consequently, odorants with different vapor pressures were diluted to correspondingly different extents in the liquid phase. Vol/vol dilutions in mineral oil to 1.0 Pa are shown for the data contributing to **Figures 3C,D**; dilutions for **Figures 3A,B** are presented in Cleland et al. (2009). Odorants used as conditioning odorants are denoted in boldface; for **Figure 3C**, the conditioning odorant was a mixture of the two boldface odorants in each odor set (with 50:50 or variable mixture ratios of the two diluted odorants as described in that study).*

diluted in mineral oil to concentrations theoretically emitting vapor-phase partial pressures of 1.0, 0.1, or 0.01 Pa as indicated. Solvent surface effects and other non-linearities were neglected. These dilutions should be considered a reduction in the variance of odor concentrations rather than true gas-phase concentration matching as could be achieved by gas chromatographic measurements. Odorants were diluted at least 18 hours in advance of each experiment to ensure an even distribution of odorant within the mineral oil solvent. These procedures have been utilized in previous studies (Cleland et al., 2002, 2009; Cleland and Narla, 2003).

ACKNOWLEDGMENTS

The authors are grateful to Kafui Nutakor and Matthew A. Haber for assistance with behavioral experiments. This research was supported by NIDCD grant DC012249.

REFERENCES

Araneda, R. C., Kini, A. D., and Firestein, S. (2000). The molecular receptive range of an odorant receptor. *Nat. Neurosci.* 3, 1248–1255.

Aungst, J. L., Heyward, P. M., Puche, A. C., Karnup, S. V., Hayar, A., Szabo, G., and Shipley, M. T. (2003). Centre-surround inhibition among olfactory bulb glomeruli. *Nature* 426, 623–629.

Baker, H. (1990). Unilateral, neonatal olfactory deprivation alters tyrosine hydroxylase expression but not aromatic amino acid decarboxylase or GABA immunoreactivity. *Neuroscience* 36, 761–771.

Baker, H., Morel, K., Stone, D. M., and Maruniak, J. A. (1993). Adult naris closure profoundly reduces tyrosine hydroxylase expression in mouse olfactory bulb. *Brain Res.* 614, 109–116.

Belluscio, L., and Katz, L. C. (2001). Symmetry, stereotypy, and topography of odorant representations in mouse olfactory bulbs. *J. Neurosci.* 21, 2113–2122.

Belluscio, L., Lodovichi, C., Feinstein, P., Mombaerts, P., and Katz, L. C. (2002). Odorant receptors instruct functional circuitry in the mouse olfactory bulb. *Nature* 419, 296–300.

Berkowicz, D. A., and Trombley, P. Q. (2000). Dopaminergic modulation at the olfactory nerve synapse. *Brain Res.* 855, 90–99.

Bhagavan, S., and Smith, B. H. (1997). Olfactory conditioning in the honey bee, Apis mellifera: effects of odor intensity. *Physiol. Behav.* 61, 107–117.

Brunjes, P. C., Smith-Crafts, L. K., and McCarty, R. (1985). Unilateral odor deprivation: effects on the development of olfactory bulb catecholamines and behavior. *Brain Res.* 354, 1–6.

Carey, R. M., Verhagen, J. V., Wesson, D. W., Pirez, N., and Wachowiak, M. (2009). Temporal structure of receptor neuron input to the olfactory bulb imaged in behaving rats. *J. Neurophysiol.* 101, 1073–1088.

Chalansonnet, M., and Chaput, M. A. (1998). Olfactory bulb output cell temporal response patterns to increasing odor concentrations in freely breathing rats. *Chem. Senses* 23, 1–9.

Cho, J. Y., Min, N., Franzen, L., and Baker, H. (1996). Rapid down-regulation of tyrosine hydroxylase expression in the olfactory bulb of naris-occluded adult rats. *J. Comp. Neurol.* 369, 264–276.

Choudhary, A. F. (2009). *Olfactory Perceptual Invariance in the Honeybee: A Psychophysical Approach*. PhD dissertation, Newcastle University, Newcastle-upon-Tyne, UK.

Cleland, T. A. (2008). "The construction of olfactory representations," in *Mechanisms of Information Processing in the Brain: Encoding of Information in Neural Populations,* eds C. Holscher and M. Munk (Cambridge, UK: Cambridge University Press), 247–280.

Cleland, T. A. (2010). Early transformations in odor representation. *Trends Neurosci.* 33, 130–139.

Cleland, T. A., Johnson, B. A., Leon, M., and Linster, C. (2007). Relational representation in the olfactory system. *Proc. Natl. Acad. Sci. U.S.A.* 104, 1953–1958.

Cleland, T. A., and Linster, C. (1999). Concentration tuning mediated by spare receptor capacity in olfactory sensory neurons: a theoretical study. *Neural. Comput.* 11, 1673–1690.

Cleland, T. A., and Linster, C. (2002). How synchronization properties among second-order sensory neurons can mediate stimulus salience. *Behav. Neurosci.* 116, 212–221.

Cleland, T. A., Morse, A., Yue, E. L., and Linster, C. (2002). Behavioral models of odor similarity. *Behav. Neurosci.* 116, 222–231.

Cleland, T. A., and Narla, V. A. (2003). Intensity modulation of olfactory acuity. *Behav. Neurosci.* 117, 1434–1440.

Cleland, T. A., Narla, V. A., and Boudadi, K. (2009). Multiple learning parameters differentially regulate olfactory generalization. *Behav. Neurosci.* 123, 26–35.

Cleland, T. A., and Sethupathy, P. (2006). Non-topographical contrast enhancement in the olfactory bulb. *BMC Neurosci.* 7, 7.

Daly, K. C., Chandra, S., Durtschi, M. L., and Smith, B. H. (2001). The generalization of an olfactory-based conditioned response reveals unique but overlapping odour representations in the moth Manduca sexta. *J. Exp. Biol.* 204, 3085–3095.

Davila, N. G., Blakemore, L. J., and Trombley, P. Q. (2003). Dopamine modulates synaptic transmission between rat olfactory bulb neurons in culture. *J. Neurophysiol.* 90, 395–404.

Duchamp-Viret, P., Duchamp, A., and Sicard, G. (1990). Olfactory discrimination over a wide concentration range. Comparison of receptor cell and bulb neuron abilities. *Brain Res.* 517, 256–262.

Ennis, M., Zhou, F. M., Ciombor, K. J., Aroniadou-Anderjaska, V., Hayar, A., Borrelli, E., Zimmer, L. A., Margolis, F., and Shipley, M. T. (2001). Dopamine D2 receptor-mediated presynaptic inhibition of olfactory nerve terminals. *J. Neurophysiol.* 86, 2986–2997.

Escanilla, O., Yuhas, C., Marzan, D., and Linster, C. (2009). Dopaminergic modulation of olfactory bulb processing affects odor discrimination learning in rats. *Behav. Neurosci.* 123, 828–833.

Fantana, A. L., Soucy, E. R., and Meister, M. (2008). Rat olfactory bulb mitral cells receive sparse glomerular inputs. *Neuron* 59, 802–814.

Fernandez, P. C., Locatelli, F. F., Person-Rennell, N., Deleo, G., and Smith, B. H. (2009). Associative conditioning tunes transient dynamics of early olfactory processing. *J. Neurosci.* 29, 10191–10202.

Firestein, S., Picco, C., and Menini, A. (1993). The relation between stimulus and response in olfactory receptor cells of the tiger salamander. *J. Physiol.* 468, 1–10.

Fletcher, M. L., Masurkar, A. V., Xing, J., Imamura, F., Xiong, W., Nagayama, S., Mutoh, H., Greer, C. A., Knopfel, T., and Chen, W. R. (2009). Optical imaging of postsynaptic odor representation in the glomerular layer of the mouse olfactory bulb. *J. Neurophysiol.* 102, 817–830.

Friedrich, R. W., and Korsching, S. I. (1997). Combinatorial and chemotopic odorant coding in the zebrafish olfactory bulb visualized by optical imaging. *Neuron* 18, 737–752.

Gall, C. M., Hendry, S. H., Seroogy, K. B., Jones, E. G., and Haycock, J. W. (1987). Evidence for coexistence of GABA and dopamine in neurons of the rat olfactory bulb. *J. Comp. Neurol.* 266, 307–318.

Gross-Isseroff, R., and Lancet, D. (1988). Concentration-dependent changes of perceived odor quality. *Chem. Senses* 13, 191–204.

Halasz, N., Johansson, O., Hokfelt, T., Ljungdahl, A., and Goldstein, M. (1981). Immunohistochemical identification of two types of dopamine neuron in the rat olfactory bulb as seen by serial sectioning. *J. Neurocytol.* 10, 251–259.

Hamilton, K. A., and Kauer, J. S. (1989). Patterns of intracellular potentials in salamander mitral/tufted cells in response to odor stimulation. *J. Neurophysiol.* 62, 609–625.

Hsia, A. Y., Vincent, J. D., and Lledo, P. M. (1999). Dopamine depresses synaptic inputs into the olfactory bulb. *J. Neurophysiol.* 82, 1082–1085.

Johnson, B. A., and Leon, M. (2000). Modular representations of odorants in the glomerular layer of the rat olfactory bulb and the effects of stimulus concentration. *J. Comp. Neurol.* 422, 496–509.

Koehl, M. A. (2006). The fluid mechanics of arthropod sniffing in turbulent odor plumes. *Chem. Senses* 31, 93–105.

Kosaka, K., Aika, Y., Toida, K., and Kosaka, T. (2001). Structure of intraglomerular dendritic tufts of mitral cells and their contacts with olfactory nerve terminals and calbindin-immunoreactive type 2 periglomerular neurons. *J. Comp. Neurol.* 440, 219–235.

Leon, M., and Johnson, B. A. (2003). Olfactory coding in the mammalian olfactory bulb. *Brain Res. Brain Res. Rev.* 42, 23–32.

Linster, C., and Cleland, T. A. (2002). Cholinergic modulation of sensory

representations in the olfactory bulb. *Neural Netw.* 15, 709–717.
Linster, C., and Hasselmo, M. E. (1999). Behavioral responses to aliphatic aldehydes can be predicted from known electrophysiological responses of mitral cells in the olfactory bulb. *Physiol. Behav.* 66, 497–502.
Luo, F., Wang, Q., Farid, N., Liu, X., and Yan, J. (2009). Three-dimensional tonotopic organization of the C57 mouse cochlear nucleus. *Hear. Res.* 257, 75–82.
Mandairon, N., Ferretti, C. J., Stack, C. M., Rubin, D. B., Cleland, T. A., and Linster, C. (2006). Cholinergic modulation in the olfactory bulb influences spontaneous olfactory discrimination in adult rats. *Eur. J. Neurosci.* 24, 3234–3244.
Mandairon, N., and Linster, C. (2009). Odor perception and olfactory bulb plasticity in adult mammals. *J. Neurophysiol.* 101, 2204–2209.
Mombaerts, P. (2004). Genes and ligands for odorant, vomeronasal and taste receptors. *Nat. Rev. Neurosci.* 5, 263–278.
Nagayama, S., Takahashi, Y. K., Yoshihara, Y., and Mori, K. (2004). Mitral and tufted cells differ in the decoding manner of odor maps in the rat olfactory bulb. *J. Neurophysiol.* 91, 2532–2540.
Nawroth, J. C., Greer, C. A., Chen, W. R., Laughlin, S. B., and Shepherd, G. M. (2007). An energy budget for the olfactory glomerulus. *J. Neurosci.* 27, 9790–9800.
Nickell, W. T., Norman, A. B., Wyatt, L. M., and Shipley, M. T. (1991). Olfactory bulb DA receptors may be located on terminals of the olfactory nerve. *Neuroreport* 2, 9–12.
Pelz, C., Gerber, B., and Menzel, R. (1997). Odorant intensity as a determinant for olfactory conditioning in honeybees: roles in discrimination, overshadowing and memory consolidation. *J. Exp. Biol.* 200, 837–847.
Pirez, N., and Wachowiak, M. (2008). *In vivo* modulation of sensory input to the olfactory bulb by tonic and activity-dependent presynaptic inhibition of receptor neurons. *J. Neurosci.* 28, 6360–6371.
Rospars, J. P., Lansky, P., Duchamp-Viret, P., and Duchamp, A. (2000). Spiking frequency versus odorant concentration in olfactory receptor neurons. *Biosystems* 58, 133–141.
Schoenfeld, T. A., and Cleland, T. A. (2005). The anatomical logic of smell. *Trends Neurosci.* 28, 620–627.
Schoenfeld, T. A., and Cleland, T. A. (2006). Anatomical contributions to odorant sampling and representation in rodents: zoning in on sniffing behavior. *Chem. Senses* 31, 131–144.
Shao, Z., Puche, A. C., Kiyokage, E., Szabo, G., and Shipley, M. T. (2009). Two GABAergic intraglomerular circuits differentially regulate tonic and phasic presynaptic inhibition of olfactory nerve terminals. *J. Neurophysiol.* 101, 1988–2001.
Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science* 237, 1317–1323.
Shusterman, R., Smear, M. C., Koulakov, A. A., and Rinberg, D. (2011). Precise olfactory responses tile the sniff cycle. *Nat. Neurosci.* 14, 1039–1044.
Stone, D. M., Wessel, T., Joh, T. H., and Baker, H. (1990). Decrease in tyrosine hydroxylase, but not aromatic L-amino acid decarboxylase, messenger RNA in rat olfactory bulb following neonatal, unilateral odor deprivation. *Brain Res. Mol. Brain Res.* 8, 291–300.
Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.
Teghtsoonian, R., Teghtsoonian, M., Berglund, B., and Berglund, U. (1978). Invariance of odor strength with sniff vigor: an olfactory analogue to size constancy. *J. Exp. Psychol. Hum. Percept. Perform.* 4, 144–152.
Tenenbaum, J. B., and Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.* 24, 629–640; discussion 652–791.
Toida, K., Kosaka, K., Aika, Y., and Kosaka, T. (2000). Chemically defined neuron groups and their subpopulations in the glomerular layer of the rat main olfactory bulb—IV. Intraglomerular synapses of tyrosine hydroxylase-immunoreactive neurons. *Neuroscience* 101, 11–17.
Treloar, H. B., Feinstein, P., Mombaerts, P., and Greer, C. A. (2002). Specificity of glomerular targeting by olfactory sensory axons. *J. Neurosci.* 22, 2469–2477.
Uchida, N., and Mainen, Z. F. (2007). Odor concentration invariance by chemical ratio coding. *Front. Syst. Neurosci.* 1:3. doi: 10.3389/neuro.06.003.2007
Verhagen, J. V., Wesson, D. W., Netoff, T. I., White, J. A., and Wachowiak, M. (2007). Sniffing controls an adaptive filter of sensory input to the olfactory bulb. *Nat. Neurosci.* 10, 631–639.
Wachowiak, M., Cohen, L. B., and Zochowski, M. R. (2002). Distributed and concentration-invariant spatial representations of odorants by receptor neuron input to the turtle olfactory bulb. *J. Neurophysiol.* 87, 1035–1045.
Wei, C. J., Linster, C., and Cleland, T. A. (2006). Dopamine D(2) receptor activation modulates perceived odor intensity. *Behav. Neurosci.* 120, 393–400.
Wellis, D. P., Scott, J. W., and Harrison, T. A. (1989). Discrimination among odorants by single neurons of the rat olfactory bulb. *J. Neurophysiol.* 61, 1161–1177.
Wesson, D. W., Varga-Wesson, A. G., Borkowski, A. H., and Wilson, D. A. (2011). Respiratory and sniffing behaviors throughout adulthood and aging in mice. *Behav. Brain Res.* 223, 99–106.
Wilson, D. A., Fletcher, M. L., and Sullivan, R. M. (2004). Acetylcholine and olfactory perceptual learning. *Learn. Mem.* 11, 28–34.
Wilson, D. A., and Stevenson, R. J. (2006). *Learning to Smell: Olfactory Perception from Neurobiology to Behavior*. Baltimore, MD: Johns Hopkins University Press.
Wilson, D. A., and Sullivan, R. M. (1995). The D2 antagonist spiperone mimics the effects of olfactory deprivation on mitral/tufted cell odor response patterns. *J. Neurosci.* 15, 5574–5581.
Wright, G. A., Choudhary, A. F., and Bentley, M. A. (2009). Reward quality influences the development of learned olfactory biases in honeybees. *Proc. Biol. Sci.* 276, 2597–2604.
Wright, G. A., Thomson, M. G., and Smith, B. H. (2005). Odour concentration affects odour identity in honeybees. *Proc. Biol. Sci.* 272, 2417–2422.
Zhu, B. T. (1993). The competitive and noncompetitive antagonism of receptor-mediated drug actions in the presence of spare receptors. *J. Pharmacol. Toxicol. Methods* 29, 85–91.

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.

Received: 03 October 2011; paper pending published: 16 November 2011; accepted: 19 December 2011; published online: 05 January 2012.
*Citation: Cleland TA, Chen ST, Hozer KW, Ukatu HN, Wong KJ and Zheng F (2012) Sequential mechanisms underlying concentration invariance in biological olfaction. Front. Neuroeng. **4**:21. doi: 10.3389/fneng.2011.00021*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 03 January 2012
doi: 10.3389/fneng.2011.00018

Glomerular latency coding in artificial olfaction

Jaber Al Yamani*[1], *Farid Boussaid*[1], *Amine Bermak*[2] *and Dominique Martinez*[3]

[1] *The University of Western Australia, Crawley, WA, Australia*
[2] *Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong, China*
[3] *UMR 7503, Laboratoire Lorrain de Recherche en Informatique et ses Applications (LORIA), CNRS, France*

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Agustín Gutierrez-Galvez, Universitat de Barcelona, Spain
Malin Sandström, International Neuroinformatics Coordinating Facility, Sweden

***Correspondence:**
Dominique Martinez, Unité Mixte de Recherche 7503, Laboratoire Lorrain de Recherche en Informatique et ses Applications, CNRS, 54506 Vandoeuvre-lès-Nancy, France.
e-mail: dominique.martinez@loria.fr

Sensory perception results from the way sensory information is subsequently transformed in the brain. Olfaction is a typical example in which odor representations undergo considerable changes as they pass from olfactory receptor neurons (ORNs) to second-order neurons. First, many ORNs expressing the same receptor protein yet presenting heterogeneous dose–response properties converge onto individually identifiable glomeruli. Second, onset latency of glomerular activation is believed to play a role in encoding odor quality and quantity in the context of fast information processing. Taking inspiration from the olfactory pathway, we designed a simple yet robust glomerular latency coding scheme for processing gas sensor data. The proposed bio-inspired approach was evaluated using an in-house SnO_2 sensor array. Glomerular convergence was achieved by noting the possible analogy between receptor protein expressed in ORNs and metal catalyst used across the fabricated gas sensor array. Ion implantation was another technique used to account both for sensor heterogeneity and enhanced sensitivity. The response of the gas sensor array was mapped into glomerular latency patterns, whose rank order is concentration-invariant. Gas recognition was achieved by simply looking for a "match" within a library of spatio-temporal spike fingerprints. Because of its simplicity, this approach enables the integration of sensing and processing onto a single-chip.

Keywords: glomerular convergence, latency coding, olfaction, electronic nose, chemical sensing, gas sensors, neuromorphic engineering

INTRODUCTION

Identifying and localizing volatile compounds is an important challenge in numerous applications such as, among others, explosive detection (civil or military security), aliment quality control, air pollution analysis (Gardner and Bartlett, 1999; Röck et al., 2008). During the past few years, there has been an increasing effort for developing low-cost microelectronic gas sensors allowing quantitative and qualitative analyses (Röck et al., 2008). Nevertheless, whatever the technology is (e.g., metal oxide, conducting polymer, piezo-electric quartz), a common feature of the sensors is their lack of selectivity (Röck et al., 2008). Gas sensors react to a large number of chemical compounds and it seems likely that the situation will continue in the long run. We can even ask whether it is useful to have selective sensors as it is cost intensive and time consuming to develop a specific sensor for every odorant to be detected. Note that lack of selectivity is also encountered in olfactory receptor neurons (ORNs). Yet, biological systems are not less very efficient. Mice and bees for example recognize learned odors in less than 200 ms (Abraham et al., 2004; Buck, 2005; Bhandawat et al., 2010; Chen et al., 2011), thereby indicating a rapid processing of the olfactory input. More strickling is the fact that primary olfactory centers in insects and vertebrates share a common design both at anatomical (cellular organization) and functional (mechanisms for olfactory coding) levels (Hildebrand and Shepherd, 1997; Strausfeld and Hildebrand, 1999; Jacquin-Joly and Lucas, 2005). The efficiency of biological olfactory systems and their similarities across species imply that nature has found an optimal solution for encoding odors (Ache and Young, 2005). The current knowledge about architectural and computational strategies used in biological systems should therefore guide the development of algorithms for processing gas sensor data.

The concept of electronic nose as a bionic system of artificial olfaction appeared about 30 years ago (Persaud and Dodd, 1982). It consisted in using an array of different sensors whose activation pattern contains the signature of the odorant to be identified, in the same way as our olfactory perception results from the activation of multiple ORNs by the odorant molecules. The analogy stops here however. The current artificial noses are only a weak imitation of the sense of smell with performance far below that of their biological counterparts. The reason is twofold.

- The organization of both systems, artificial and biological, is not comparable. In electronic noses, the sensing elements are mapped on a two-dimensional array without particular organization. In biological olfactory systems on the contrary, the sensory neurons distributed on insect antennae or in vertebrate epithelia converge onto individually spherical neuropils, called glomeruli, where all synaptic connections between ORNs and second-order neurons are made (**Figure 1**). Olfactory glomeruli are individually identifiable across animals and are functionally specialized in terms of odor processing (Rospars, 1988; Baier and Korsching, 1994). It is now well established that single ORN projects onto a few (most often only one) glomeruli and that all ORNs from a given glomerulus express the same olfactory

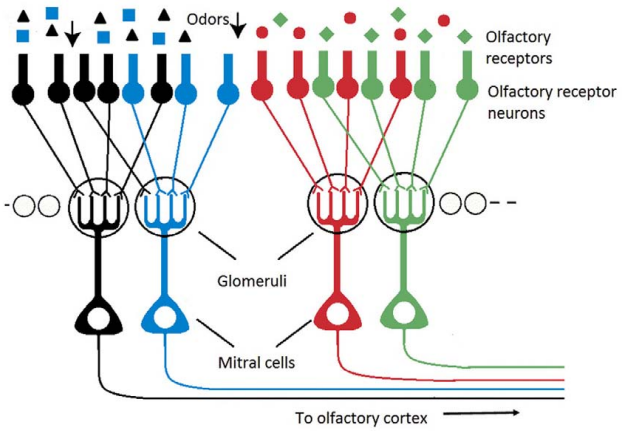

FIGURE 1 | Schematic of the olfactory pathway (olfactory bulb) illustrating that the ORNs expressing the same OR converge onto spherical individual neuropils, called glomeruli, where all synaptic connections between ORNs and second-order neurons (mitral cells) are made.

receptor (OR), (Mombaerts, 1996, 2001; Vosshall et al., 2000). The ORNs of a given glomerulus are thus homogeneous in terms of receptor type. It has been shown however that in an homogeneous population with same OR, the ORNs present variable odorant response properties (Grosmaitre et al., 2006; Grémiaux et al., 2011). Computational advantages for such a variability are still unclear but they could include lower detection thresholds and wider dynamic ranges by averaging many independent inputs (Grosmaitre et al., 2006).

- The processing of gas sensor arrays currently based on statistical methods (e.g., discriminant analysis, kernel methods) is very far from the biological reality (neurons firing action potentials or spikes). Second-order neurons (mitral cells in the vertebrate olfactory bulb, projection neurons in the insect antennal lobe) were found to encode information in a few spikes, using relative latencies after stimulus onset as an information carrier (Hopfield, 1995; Margrie and Schaefer, 2003; Junek et al., 2010; Belmabrouk et al., 2011; Smear et al., 2011). Timing of activity relative to the sampling behavior (e.g., sniffing in rodents) leads to a concentration-invariant code for odor identity (Hopfield, 1995; Margrie and Schaefer, 2003) and may play a significant role in the context of fast information processing, a rat for example recognizing a novel odor in less than 200 ms (Abraham et al., 2004; Wesson et al., 2008) and a drosophila in 90 ms (Bhandawat et al., 2010).

Taking inspiration from the organizational and functional characteristics of the olfactory pathway (glomerular convergence and latency coding), we designed a simple yet robust glomerular latency coding scheme for processing gas sensor data.

MATERIALS AND METHODS

SENSOR ARRAY FABRICATION AND CHARACTERIZATION

Using an in-house 5 μm, 2-metal, 1-poly process, we have fabricated a 4×4 gas sensor array consisting of 16 SnO_2 sensors, corresponding to different combinations of dopants and catalysts. The structure of the fabricated 4×4 gas sensor array is shown in **Figure 2**. Sensors belonging to the same column share the same metal catalyst (Pt, Ag, or Au), while row wise sensors share the same dopant (B, P, or H). Note that a single row of sensors has no implanted dopant (ND) and a single column of sensors has no metal catalyst (NC). Each individual sensor integrates a microhotplate (MHP) heater (**Figure 2**), which consists of a membrane stretched over a rigid frame (Graf, 2007). The membrane plays two roles: (i) it acts as a support for the functional elements of the sensor (e.g., heater and sensing layer), and (ii) it provides thermal isolation of the hot sensing area through its low heat conductivity (Graf, 2007). The presence of a target gas is detected through changes in the conductance of the sensing film, which depends on the gas type and concentration, its operating temperature but also on the deposited sensing film. In our array, a 100-nm SnO_2 film was deposited on the MHP structure, using sputtering and lift-off techniques. Dopant implantation (B, P, or H) was selectively carried out at a dose of 5×10^{12} cm^{-2} and energy of 30 keV. Different catalysts (Pt, Ag, or Au) were subsequently deposited on individual sensing films by RF sputtering. Each MHP is a 190×190-μm^2 oxide/low-stress nitride/oxide (O/N/O) multilayer membrane. A 2.8-μm air gap separates MHP and substrate to reduce heat losses to the substrate. The air gap was formed by etching a sacrificial polysilicon layer. A serpentine Pt microheater was patterned at the center of the MHP, using sputtering and lift-off techniques (**Figure 2**). The microheater has a width of 10 μm and a thickness of 100 nm. Its resistance is about 297 Ω at room temperature and 441 Ω at the 300°C operating temperature. Plasma enhanced chemical vapor deposition (PECVD) was used to deposit an insulating oxide/nitride/oxide multilayer on top of the microheater. The resistance change of each individual sensing film is measured across two Pt electrodes. In the fabricated 4×4 gas sensor array, the catalysts play the role of the receptor and the dopants add heterogeneity in the sensors' responses.

The fabricated electronic nose was characterized in a controlled laboratory environment, using an automated gas delivery system (**Figure 3**). Test gases used in the experiments were methane, carbon monoxide, and ethanol. Flow rates were set by adjusting the voltage of computer controlled mass-flow-controllers (MFCs). The control of the gas concentration, within the 2.5-cm radius 30 cm^3 cylinder testing chamber, was achieved by mixing the target gas with dry air at different flow rates. The relative small size of the chamber contributes to the uniformity of the gas concentration across the sensor array. Prior to each gas exposure, the surface of the sensor array was cleaned by injecting dry air. The sensor array was subsequently exposed to the target gas for up to 300 s. Throughout each clean-and-expose-cycle, resistance variations across the array were recorded simultaneously using 10-bit digital multimeters. The sensor steady state resistance was sampled just before the end of the analyte injection period.

SPIKE LATENCY CODING

In the generalist olfactory pathway, ORNs are sensitive to multiple odorants, and each odorant activates thousands of ORNs, leading to a massive combinatorial code at the receptor level. Yet this code is not only combinatorial, it is also structured temporally by the

Sensor number	Sensor type
1	Au/P
2	Ag/P
3	Pt/P
4	NC/P
5	Au/B
6	Ag/ B
7	Pt/ B
8	NC/ B

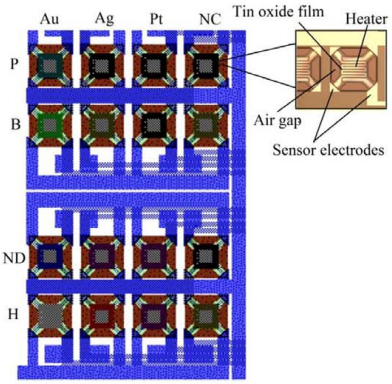

Sensor number	Sensor type
9	Au/ ND
10	Ag/ ND
11	Pt/ ND
12	NC/ ND
13	Au/H
14	Ag/ H
15	Pt/ H
16	NC/ H

FIGURE 2 | The fabricated 4 x 4 SnO_2 gas sensor array with the structure of an individual sensor highlighted on the right. There are 16 SnO_2 sensors, corresponding to different combinations of dopants and catalysts. Sensors belonging to the same column share the same metal catalyst (Pt, Ag, or Au), while row wise sensors share the same dopant (B, P, or H). A single row has no dopant (ND) and a single column has no metal catalyst (NC). Note that the above sensor numbering is kept for all subsequent figures.

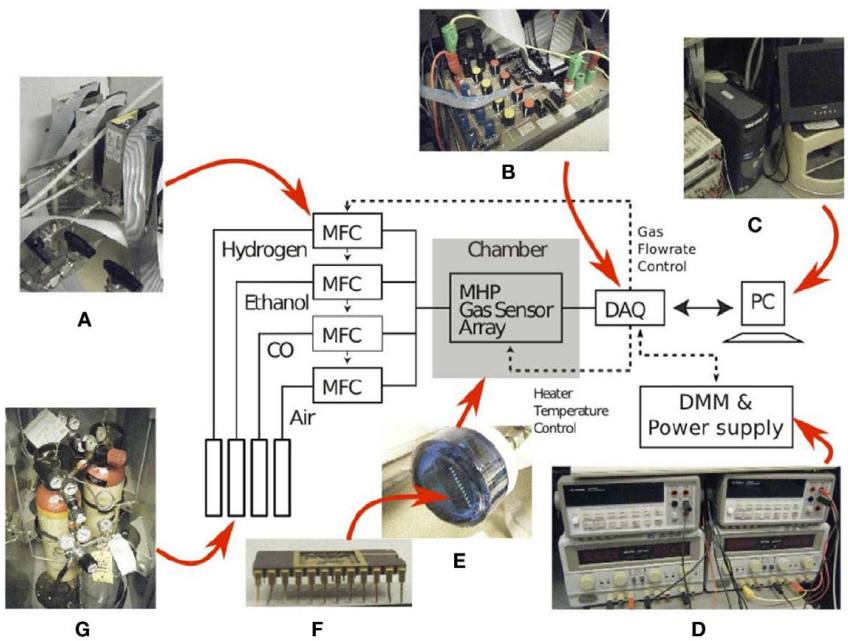

FIGURE 3 | Experimental setup used to characterize the gas sensor array: (A) four mass-flow-controllers, (B) a data-acquisition-board (DAQ), (C) a desktop computer, (D) power supplies and digital multimeters, (E) a chamber housing the gas sensor array packaged in a (F) 24-pins Ceramic Dual-Inline-Package (CDIP24) and (G) three gases namely H_2, ethanol, and CO.

sampling behavior of the animal (e.g., sniffing in rodents) which provides a temporal frame of reference. Recent evidence shows that mice are able to discriminate between olfactory inputs merely on the basis of timing information relative to the respiration cycle (Smear et al., 2011). Electrophysiological recordings in the mouse olfactory bulb have revealed that the firing latency of mitral cells relative to the respiration cycle depends on odor intensity in a logarithmic way (Margrie and Schaefer, 2003) and theoretical works have suggested that such a logarithmic transformation makes relative latencies invariant to odor concentration (Hopfield, 1995).

We have exploited the above ideas to design a spike latency coding scheme that converts the sensor responses into a unique sequence of latency spikes (Martinez et al., 2006; Ng et al., 2009; Chen et al., 2011). The firing latency t_i associated to the i-th sensor exposed to a target gas j is proportional to the logarithm of the sensor resistance

$$t_i = \frac{\ln R_{ij}}{\gamma_{ij}} \quad (1)$$

where γ_{ij} is a parameter which depends both on the type of gas j and sensing material i and R_{ij} is the resistance of sensor i when exposed to target gas j. The resistance R_{ij} is here modeled as a power law (Gardner and Bartlett, 1999; Yamazoe and Shimanoe, 2008) leading to

$$t_i = \frac{\ln \alpha_{ij}}{\gamma_{ij}} - \ln C_j \text{ given that } R_{ij} = \alpha_{ij} C_j^{-\gamma_{ij}} \quad (2)$$

where α_{ij} is a parameter depending both on sensor i and gas j and C_j is the concentration of gas j. Considering the previous equation, the relative latency between sensors p and q can be expressed as

$$t_{pq} = \frac{\ln \alpha_{pj}}{\gamma_{pj}} - \frac{\ln \alpha_{qj}}{\gamma_{qj}} \quad (3)$$

Note that the relative latency is concentration independent (C_j does not appear in Eq. 3) while being gas dependent (both α and γ depend on the type of gas). As a result, the rank order (i.e., firing order of sensors) can be used as a signature (**Figure 4**, top) and gas recognition can be simply achieved by looking for a match within a library of spatio-temporal spike signatures.

GLOMERULAR LATENCY CODING

In the olfactory pathway, each glomerulus receives inputs from ORNs expressing the same OR (**Figure 1**). We have emulated glomerular convergence in the fabricated gas sensor array, by making an analogy between OR expressed and metal catalyst used. In this way, we defined four glomeruli (considering the four types of sensors Pt, Ag, Au, and NC), each one receiving four sensory inputs. We further considered that the elementary unit of the olfactory code is a glomerular module that consists of a glomerulus together with its associated second-order neurons (Mori et al., 1999; Koulakov et al., 2007). The onset latency of a glomerular module p was defined as

$$t_p = \frac{\sum_{i \in p} \ln\left(R_{ij}\right)}{\sum_{i \in p} \gamma_{ij}} \quad (4)$$

where the sensors i are of the same type (i.e., with same catalyst p). Considering a power law sensor response leads to

$$t_p = \frac{\sum_{i \in p} \ln\left(\alpha_{ij}\right)}{\sum_{i \in p} \gamma_{ij}} - \ln C_j \quad (5)$$

and the relative latency between two glomerular modules p and q (catalysts p and q) is expressed as

$$t_{pq} = \frac{\sum_{i \in p} \ln\left(\alpha_{ij}\right)}{\sum_{i \in p} \gamma_{ij}} - \frac{\sum_{i \in q} \ln\left(\alpha_{ij}\right)}{\sum_{i \in q} \gamma_{ij}} \quad (6)$$

As for the spike latency coding, the relative latency is concentration independent (C_j does not appear in Eq. 6) while being gas dependent (both α and γ depend on the type of gas). The glomerular latency code (**Figure 4**, bottom) is however more compact than the spike latency code (**Figure 4**, top). For both spike and glomerular latency codes, the parameters required to convert the sensor resistances into firing times are the sensitivities γ_{ij} (see Eqs 1 and 4). These parameters were estimated empirically by calibration, i.e., by exposing the sensor array to each target gas j delivered at different known concentrations C_j ranging from 20 to 200 ppm, and determining the slope of the linear regression of ln R_{ij} versus ln C_j. Because individual γ_{ij} parameters depend both on the type of gas j and sensing material i, 64 parameters have been extracted given that we considered four target gases and 16 different sensors. The range of sensitivity values was found to be $0 < \gamma_{ij} < 1.1$.

RESULTS

SENSORS WITH THE SAME CATALYST EXHIBIT SIMILAR BEHAVIORS

Hydrogen, methane, carbon monoxide, and ethanol, at concentrations ranging from 20 to 200 ppm, were used to characterize the fabricated gas sensor array, operated at 300°C. Examples of recorded response curves of individual sensors exposed to ethanol at different concentrations are shown in **Figure 5**. Note that following each gas exposure (at a given concentration), the gas chamber was subsequently cleaned by injecting dry air. After the cleaning process, the sensor array was exposed to the same gas but at another concentration level. Observe that increasing gas concentration has the effect of decreasing the sensor resistance, which follows a power law. This is the case because all target gases are reducing gases. When dry air is injected, the sensor resistance increases back to the initial baseline resistance. Interestingly, we found that the sensors sharing the same catalyst exhibited similar sensitivities to the target gases. As an example, the sensitivities of the sensors to ethanol, as measured by the γ_{ij} parameters, were 0.95 ± 0.09 (mean $\pm$ SD) with catalyst Pt (red curves in **Figure 5**), 0.22 ± 0.06 with Au (black curves in **Figure 5**), and 0.29 ± 0.08 with Ag (blue curves in **Figure 5**).

One major challenge to gas identification is the inherent drift of gas sensors, which results in temporal variations of the sensor response with repeated experiments. **Figure 6** shows the drift of the sensors as defined by $(R_{ij} - R'_{ij})/R_{ij}$ where R_{ij} is the resistance of sensor i exposed to gas j measured in the first cycle while R'_{ij} is the sensor resistance measured after the experiment has been repeated n times. The drift sensitivities, computed as the slopes of the linear regressions in **Figure 6**, were 0.76 ± 0.002 (mean $\pm$ SD) with catalyst Pt (red curves in **Figure 6**), 0.23 ± 0.0003 with Au (black curves in **Figure 6**), and 0.25 ± 0.001 with Ag (blue curves in **Figure 6**). As within catalyst-group variances are small, these results indicate that sensors with the same catalyst exhibit similar drift behaviors.

SPIKE LATENCY CODING GREATLY SIMPLIFIES THE TASK OF GAS RECOGNITION

Most current approaches to gas identification rely on statistical pattern-recognition techniques (Gutierrez-Osuna, 2002). The silicon area and thus cost associated to their implementation are too

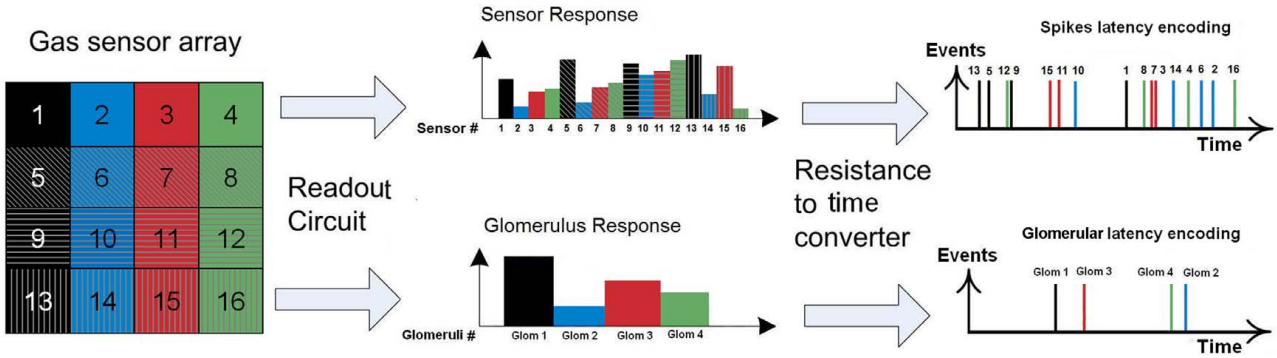

FIGURE 4 | Spike and glomerular latency coding. Top: spike latency coding (Eq. 1) converts the response of the fabricated 4×4 sensor array into a unique sequence of 16 spikes with the firing latency mapping the strength of the stimulation across the array. Sensor numbering corresponds to that of **Figure 2**, with a single color associated to sensors sharing the same catalyst (Au, Ag, Pt, or NC). Bottom: glomerular latency coding (Eq. 4) converts the sensor array into a unique sequence of four spikes which correspond to the four virtual glomeruli defined by grouping the sensors according to their catalysts.

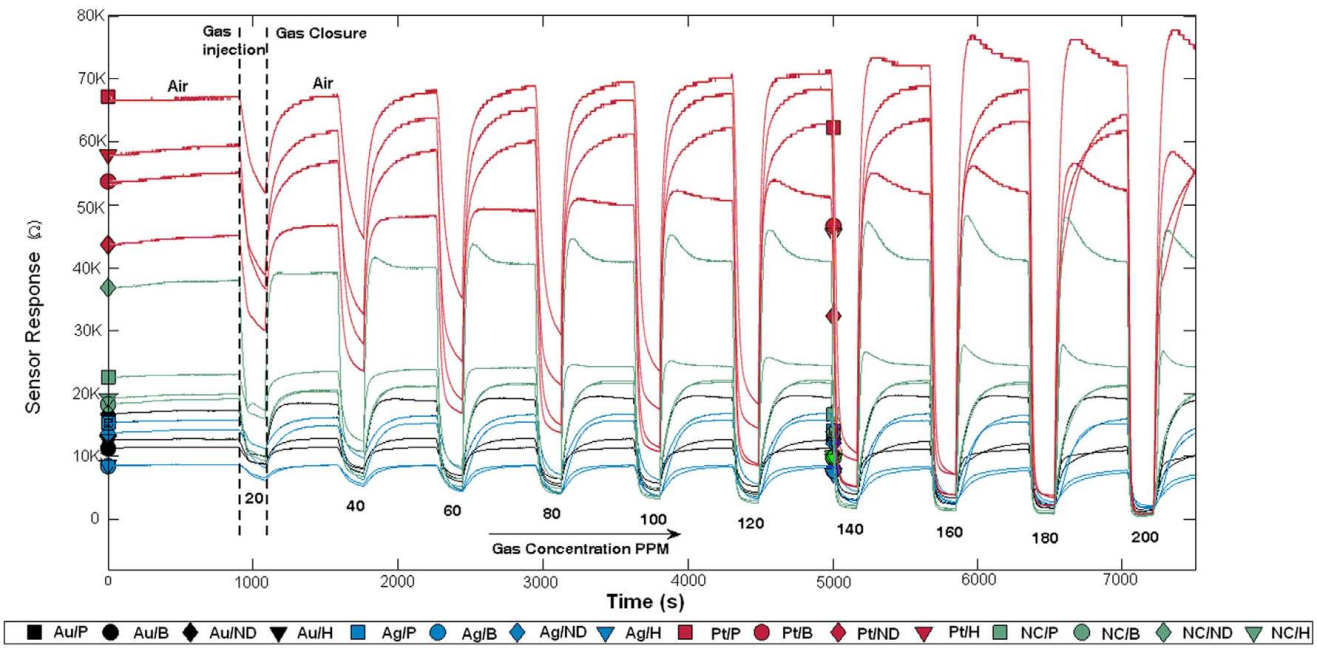

FIGURE 5 | Sensor array response to ethanol at concentrations ranging from 20 to 200 ppm. After each exposure, the sensor array is cleaned by injecting dry air. Observe that sensor resistances decrease during exposures because all target gases are reducing gases. Sensors sharing the same catalyst are plotted with the same color. ND stands for No Dopant and NC for No Catalyst.

prohibitive to envision their on-chip integration with the sensor array (Gutierrez-Osuna, 2002). To address this issue, we have developed a bio-inspired encoding scheme that can convert the response of the sensor array into a unique sequence of spikes, with the firing delay mapping the strength of the stimulation across the array. In spike latency coding (Eq. 1 in Materials and Methods, see also **Figure 4** top), the inter-spike interval is concentration independent while the firing order of the spikes is gas dependent. This means that the rank order (i.e., firing order of sensors) can be used as a signature to identify the target gas. We have experimentally validated this encoding scheme using the fabricated in-house 4×4 SnO_2 gas sensor array. For any target gas (methane, hydrogen, ethanol, and carbon monoxide) the rank order of the spike latency code did not change much when the gas concentration increased. The corresponding rank order signatures are shown in **Figure 7A**, with a correct detection rate ranging from 80 to 99.1% (**Table 1**). These results suggest that the traditionally complex and computationally intensive task of gas recognition can be dramatically simplified to the task of looking for a match within a library of spatio-temporal spike signatures. To assess the benefit of spike

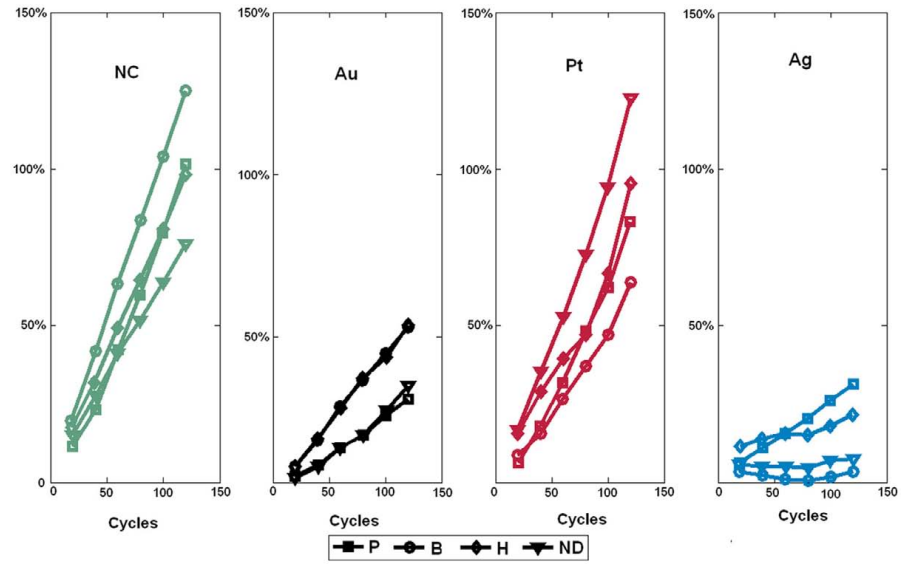

FIGURE 6 | Drift behavior of sensors sharing the same catalyst as a function of the number of clean-expose cycles. The injected gas was ethanol at 200 ppm. Note that the sensor drift behavior is mainly determined by the catalyst and not the type of dopant. ND stands for No Dopant and NC for No Catalyst.

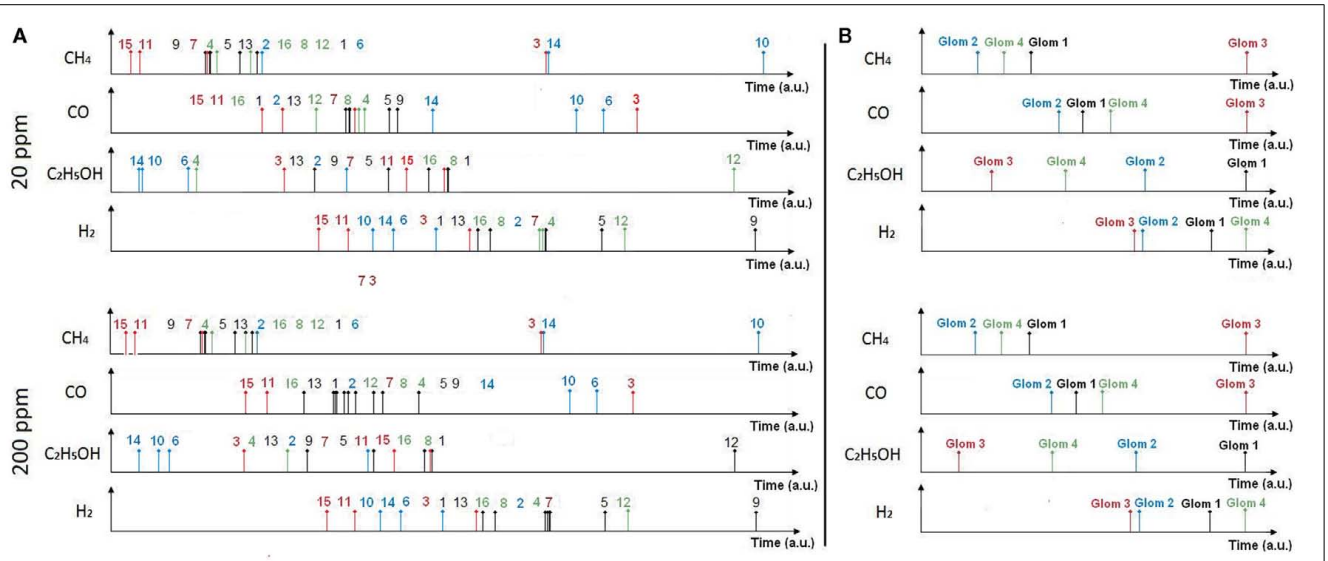

FIGURE 7 | Rank order signatures of the target gases at concentrations of 20 and 200 ppm using spike latency coding (A) and glomerular latency coding (B). In spike latency coding **(A)** the rank order corresponds to the sensor labeling given in **Figure 2**. In glomerular latency coding **(B)**, the rank order (1, 2, 3, 4) corresponds to four virtual glomeruli formed by catalysts Au, Ag, Pt, and NC, respectively. Note that the rank order signature is 100% concentration-invariant in the case of glomerular latency coding. The resistance to time conversion was achieved using Eqs 1 and 4 for spike and glomerular latency codes, respectively.

encoding in terms of pattern separability, we projected the 16-dimensional spike vectors onto a two-dimensional space by using principal component analysis (PCA). The first two eigenvectors accounted for 99.7% of the variance and four clusters corresponding to the target gases were easily identifiable (see **Figure 8B**). For comparison, the raw sensor data (sensor resistances) were also projected using PCA and a much larger within-class scatter was obtained (**Figure 8A**).

GLOMERULAR CONVERGENCE INCREASES ROBUSTNESS

In the mammalian olfactory bulb and the insect antennal lobe, glomeruli are generally thought to represent functional units of olfactory coding. A glomerulus receives axonal inputs from thousands of ORNs that all express the same OR. This glomerular convergence is believed to provide improved signal-to-noise ratio as well as increased sensitivity. We have emulated glomerular convergence using the fabricated in-house 4×4 SnO_2 gas sensor array,

Table 1 | Performance comparison between spike and glomerular latency coding.

Target gas	Gas recognition performance (%)	
	Spike latency coding (16-spike rank order; %)	Glomerular latency coding (4-spike rank order; %)
Hydrogen (H_2)	99.1	100
Carbon monoxide (CO)	80	100
Ethanol (C_2H_5OH)	91.82	100
Methane (CH_4)	99.1	100

by noting the possible analogy between OR protein expressed and metal catalyst used across the sensor array. In this way, we defined four glomeruli by considering the four types of sensors (catalysts Au, Pt, Ag, and NC), each one processing inputs from four sensors (glomerular latency coding given by Eq. 4 in Materials and Methods, see also **Figure 4** bottom). The proposed glomerular latency coding serves two computational functions. (i) It leads to a more compact odorant representation (e.g., four glomerular spikes in **Figure 7B**) than that available at the sensory level. (ii) It averages out uncorrelated noise by summation of sensory inputs of the same type, resulting in 100% accuracy for our dataset (**Table 1**). These bio-inspired computational advantages were obtained because the processed sensors had similar gas affinity, similarly to ORNs of the same type converging to a given glomerulus.

To assess the gas affinity of the sensors sharing the same catalyst, we projected the sensors' responses of each gas onto a two-dimensional space by using PCA (**Figure 9**). Four separate clusters corresponding to the four catalysts were identified. The clusters were well separated for three gases (methane, hydrogen, and carbon monoxide) and the model thus mimics the formation of individual glomeruli. For ethanol however, only the Pt cluster is well separated, with the other catalyst clusters exhibiting similar response to ethanol. This difference can be explained from the mean sensitivities to ethanol ($\gamma_{ij} = 0.95$ on average for Pt versus 0.22 and 0.29 for Au and Ag). For comparison, we performed a similar analysis by grouping the sensors according to their dopants. The PCA plots did not reveal any particular clusters for any of the target gases specified above (see **Figure 9**). All together, these results suggest that sensors sharing the same catalyst have similar gas affinity to the majority of the target gases (methane, hydrogen, and carbon monoxide) thereby validating the analogy with ORNs expressing the same OR.

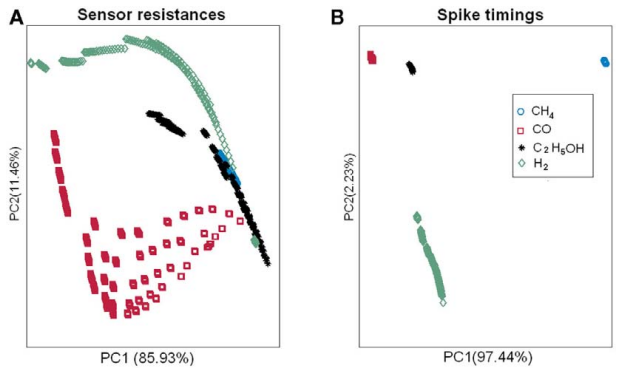

FIGURE 8 | Pattern separability achieved with spike latency coding (B) as compared to the raw sensor data (A). For spike latency coding **(B)**, the 16-dimensional spike vectors were projected onto a two-dimensional space by using principal component analysis (PCA). The first two eigenvectors accounted for 99.7% of the variance and four clusters corresponding to the target gases were easily identifiable. For comparison, the raw sensor data (sensor resistances) were also projected using PCA and a much larger within-class scatter was obtained **(A)**.

DISCUSSION

SENSORY MISMATCH

The front-end of the olfactory pathway comprises a massive number (∼10–100 million) of ORNs, each of which selectively expresses one or a few genes from a large (∼1,000) family of receptor proteins (Buck, 2005). This massively redundant representation improves signal-to-noise ratio, providing increased sensitivity in the subsequent processing layers (Pearce et al., 2003; Koickal et al., 2007). Unlike the biological olfactory system, the electronic nose uses very few sensors with commonly one replica of only several sensor types. Progress in the fabrication of largely redundant and diverse arrays has been hindered by process complexity, variability, and cost. This is not surprising given the number of process variables (e.g., sputtering power, substrate temperature, reactant concentration, etc.) that can affect a material (Pearce et al., 2003; Röck et al., 2008). Even if a high degree of sensory diversity and redundancy was implemented, the number of connections required to read-out such an array would still be unmanageable

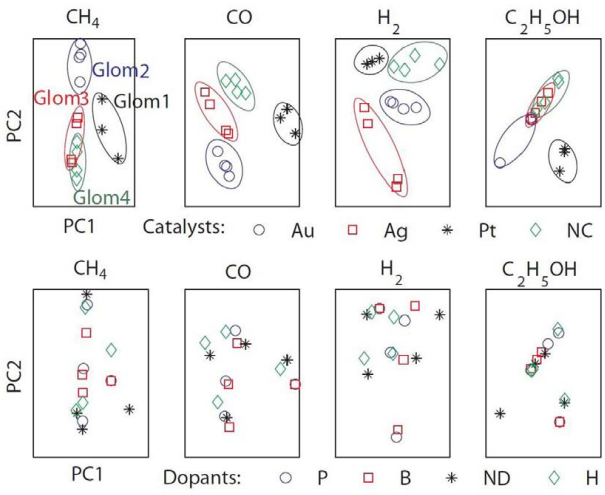

FIGURE 9 | The sensors' responses to target gases (methane, hydrogen, ethanol, and carbon monoxide) were projected onto two-dimensional spaces by using PCA. The first two eigenvectors accounted for more than 99% of the variance. Top: grouping the sensors according to their catalysts revealed four identifiable clusters mimicking the formation of biological glomeruli (Glom1, 2, 3, and 4). Bottom: grouping the sensors according to their dopants did not reveal any particular clustering.

using conventional read-out techniques. In this work, we fixed the size of our gas sensor array to 4×4 to limit the complexity and cost associated to the fabrication process.

ROLE OF THE CATALYST

When an odorant binds to an odor receptor, the receptor undergoes structural changes. This activates the olfactory-type G protein that is inside the OR neuron. Molecular studies have suggested that each ORN expresses a single type of OR genes. ORN specificity has been ascribed to that of its OR protein. Achieving analogous selective sensor response toward particular target gases involves altering the sensing surface. This can be achieved, during sensor fabrication, through the introduction of dopants (e.g., P, B, H) and/or metal additives (e.g., Pt, Ag, Au). In the fabricated 4×4 SnO_2 sensor array, we emulated ORN specificity through the use of catalysts, with each playing the role of the OR protein. This is possible because catalysts are known to dominate the sensing mechanism and the equilibrium rate of the redox reactions between detected gas and sensor surface (Vlachos et al., 1997). This may explain why sensors sharing the same catalysts are seen to exhibit the same drift behavior (**Figure 6**). Mechanisms proposed to explain the catalyst control are all based on oxygen reactions at the sensor surface. The sensitivity provided by a given metal catalyst has been shown to depend on its work function because the depletion region formed inside the semiconductor has the effect of reducing the electron availability of semiconducting grains. This makes their active size smaller than the geometrical one. Furthermore, experimental results have shown that sensitivity provided by a given catalyst is directly correlated to the active grain size (Vlachos et al., 1997). In regard to the implanted dopants, their effect is to further enhance the sensitivity of the sensor by introducing impurities into the sensing material. These impurities take up some of the substitutional and interstitial sites in the lattice, creating more oxygen vacancies for adsorption.

GLOMERULAR CONVERGENCE

Experimental studies have established that ORNs expressing the same OR converge precisely to a single glomerulus (or a small number of glomeruli). The implications of this convergence, in terms of coding are believed to be twofold: improved sensitivity so as to ensure detection and increased signal-to-noise ratios by averaging out of uncorrelated noise (Laurent, 1999). Raman et al. (2006) have sought to emulate this glomerular convergence by dynamically modulating the operating temperature so as to create thousands of virtual sensors.

However, the specificity of the sensors belonging to the same glomerulus cannot be controlled as it is in our technique with the analogy made between the receptor protein expressed in ORNs and the metal catalyst used in SnO_2 sensors. Although capable of emulating a large number of pseudo-sensors, the approach followed in Raman et al. (2006) provides highly correlated and redundant data given that only two real sensors are used. In contrast, the sensors sharing the same catalysts in our fabricated gas sensor array exhibited different dopants to mimic the variability in terms of odorant response properties found in ORNs expressing the same OR (Grosmaitre et al., 2006; Grémiaux et al., 2011).

PRACTICAL LIMITATIONS

The actual number of virtual glomeruli was limited to 4, because of process complexity and cost considerations. Increasing the number of glomeruli to n, would increase the discriminability capability with the number of possible rank order codes increasing from 4! (i.e., 24) to $n!$ Increasing the number of ORNs converging into a single glomerulus would further improve the sensitivity. Concentration-invariance of glomerular latency coding resulted from the power law sensor response. However, this relationship might break down at very low concentrations (Gardner and Bartlett, 1999). Our experimental setup does not allow us to control gas concentration at those levels, which would be required for applications such as trace detection of explosives or drugs. Another important factor that needs to be considered is the impact of long-term drift which occurs as a result of dynamic processes (e.g., poisoning or aging) or environmental changes (e.g., temperature and pressure conditions). Our experimental results (**Figure 6**) show that sensors sharing the same catalyst (i.e., belonging to the same glomerulus) exhibit the same long-term drift behavior. This observation was also reported by Sulz et al. (1993) who suggested that the overall sensor drift behavior is mainly determined by the chosen catalyst rather than the chosen dopant (Sulz et al., 1993). Vlachos et al. (1997) explained that such a behavior may be most probably due to the fact that catalysts dominate the sensing mechanism and the equilibrium rate (Vlachos et al., 1997). The above results suggest that the specificity of the catalyst and thus that of its associated glomerulus is still preserved after long-term drift. These observations could explain why the rank order of the glomerular latency code did not change after hundreds of measurements taken.

HARDWARE CONSIDERATIONS

In existing electronic noses, all sensor data need to be transferred to the pattern-recognition engine to identify ambient gases (Gutierrez-Osuna, 2002). The size and power consumption required to support this traditional architecture are prohibitive because most of the current approaches for processing multivariate sensor data are direct applications of statistical and chemometric pattern-recognition techniques (Gutierrez-Osuna, 2002). In addition, these techniques rapidly break down when the dimensionality of the input space becomes large. As a result, the co-integration of sensor and circuitry has been so far limited to on-chip signal amplification, conditioning, and/or analog-to-digital conversion (Hagleitner et al., 2001; Hierlemann and Baltes, 2003; Graf, 2007). In contrast, the proposed glomerular latency coding turns the task of gas recognition into a simple code matching task. The matching can be carried out by means of simple XOR gates, enabling the integration of sensing, and processing elements on a single-chip. Furthermore, "fingerprint matching" can easily handle large sensor arrays, since it can be aborted as soon as the detected spike does not match. The proposed scheme requires a single off-line calibration to determine gas fingerprints, while gas parameters γ_i can be stored on-chip.

Although the proposed glomerular latency coding constitutes an excellent tradeoff between classification accuracy and implementation complexity, further work is required to generalize these results to other types of sensors and target gases.

REFERENCES

Abraham, N. M., Spors, H., Carleton, A., Margrie, T. W., Kuner, T., and Schaefer, A. T. (2004). Maintaining accuracy at the expense of speed: stimulus similarity defines odor discrimination time in mice. *Neuron* 44, 865–876.

Ache, B. W., and Young, J. M. (2005). Olfaction: diverse species, conserved principles. *Neuron* 48, 417–430.

Baier, H., and Korsching, S. (1994). Olfactory glomeruli in the zebrafish form an invariant patternand are identifiable across animals. *J. Neurosci.* 14, 219–230.

Belmabrouk, H., Nowotny, T., Rospars, J.-P., and Martinez, D. (2011). Interaction of cellular and network mechanisms for efficient pheromone coding in moths. *Proc. Natl. Acad. Sci. U.S.A.* 108, 19790–19795.

Bhandawat, V., Maimon, G., Dickinson, M. H., and Wilson, R. I. (2010). Olfactory modulation of flight in *Drosophila* is sensitive, selective and rapid. *J. Exp. Biol.* 213, 3625–3635.

Buck, L. B. (2005). Unraveling the sense of smell (Nobel lecture). *Angew. Chem. Int. Ed. Engl.* 44, 6128–6140.

Chen, H. T., Ng, K. T., Bermak, A., Law, M. K., and Martinez, D. (2011). Spike latency coding in a biologically inspired micro-electronic nose. *IEEE Trans. Biomed. Circuits Syst.* 5, 160–168.

Gardner, J. W., and Bartlett, P. N. (1999). *Electronic Noses, Principles and Applications.* Oxford: Oxford University Press.

Graf, M. (2007). *CMOS Hotplate Chemical Microsensors.* Berlin: Springer Verlag.

Grémiaux, A., Nowotny, T., Martinez, D., Lucas, P., and Rospars, J.-P. (2011). Modelling the signal delivered by a population of first-order neurons in a moth olfactory system. *Brain Res.* doi:10.1016/j.brainres.2011.09.035

Grosmaitre, X., Vassalli, A., Mombaerts, P., Shepherd, G. M., and Ma, M. (2006). Odorant responses of olfactory sensory neurons expressing the odorant receptor MOR23: a patch clamp analysis in gene-targeted mice. *Proc. Natl. Acad. Sci. U.S.A.* 103, 1970–1975.

Gutierrez-Osuna, R. (2002). Pattern analysis for machine olfaction: a review. *IEEE Sens. J.* 2, 189–202.

Hagleitner, C., Hierlemann, A., Lange, D., Kummer, A., Kerness, N., Brand, O., and Baltes, H. (2001). Smart single-chip gas sensor microsystem. *Nature* 414, 293–296.

Hierlemann, A., and Baltes, H. (2003). CMOS-based chemical microsensors. *Analyst* 128, 15–28.

Hildebrand, J. G., and Shepherd, G. M. (1997). Mechanisms of olfactory discrimination: converging evidence for common principles across phyla. *Annu. Rev. Neurosci.* 20, 595–631.

Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376, 33–36.

Jacquin-Joly, E., and Lucas, P. (2005). Pheromone reception and transduction: mammals and insects illustrate converging mechanisms across phyla. *Curr. Top. Neurochem.* 4, 75–105.

Junek, S., Kludt, E., Wolf, F., and Schild, D. (2010). Olfactory coding with patterns of response latencies. *Neuron* 67, 872–884.

Koickal, T. J., Hamilton, A., Tan, S. L., Covington, J. A., Gardner, J. W., and Pearce, T. C. (2007). Analog VLSI circuit implementation of an adaptive neuromorphic olfaction chip. *IEEE Trans. Biomed. Circuits Syst.* 54, 60–73.

Koulakov, A., Gelperin, A., and Rinberg, D. (2007). Olfactory coding with all-or-nothing glomeruli. *J. Neurophysiol.* 98, 3134–73142.

Laurent, G. (1999). A systems perspective on early olfactory coding. *Science* 286, 723–728.

Margrie, T. W., and Schaefer, A. T. (2003). Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J. Physiol.* 546, 363–374.

Martinez, D., Rochel, O., and Hugues, E. (2006). A biomimetic robot for tracking specific odors in turbulent plumes. *Auton. Robots* 20, 185–195.

Mombaerts, P. (1996). Targeting olfaction. *Curr. Opin. Neurobiol.* 6, 481–486.

Mombaerts, P. (2001). How smell develops. *Nat. Neurosci.* 4, 1192–1198.

Mori, K., Nagao, H., and Yoshihara, Y. (1999). The olfactory bulb: coding and processing of odor molecule information. *Science* 286, 711–715.

Ng, K. T., Chen, H. T., Boussaid, F., Bermak, A., and Martinez, D. (2009). A robust spike-based gas identification technique for SnO2 gas sensors. *IEEE Int. Symp. Circuits Syst.* 553–556.

Pearce, T. C., Schiffman, S. S., Nagle, H. T., and Gardner, J. W. (2003). *Handbook of Machine Olfaction.* Weinheim: Wiley-VCH.

Persaud, K., and Dodd, G. H. (1982). Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature* 299, 352–355.

Raman, B., Sun, P. A., Gutierrez-Galvez, A., and Gutierrez-Osuna, R. (2006). Processing of chemical sensor arrays with a biologically inspired model of olfactory coding. *IEEE Trans. Neural Netw.* 17, 1015–1024.

Röck, F., Barsan, N., and Weimar, U. (2008). Electronic nose: current status and future trends. *Chem. Rev.* 108, 705–725.

Rospars, J. P. (1988). Structure and development of the insect antennodeutocerebral system. *Int. J. Insect Morphol. Embryol.* 17, 243–294.

Smear, M., Shusterman, R., O'Connor, R., Bozza, T., and Rinberg, D. (2011). Perception of sniff phase in mouse olfaction. *Nature* 479, 397–400.

Strausfeld, N. J., and Hildebrand, J. G. (1999). Olfactory systems: common design, uncommon origins? *Curr. Opin. Neurobiol.* 9, 634–639.

Sulz, G., Kühner, G., Reiter, H., Uptmoor, G., Schweizer, W., Low, H., Lacher, M., and Steiner, K. (1993). Ni, In and Sb implanted Pt and V catalysed thin-film SnO_2 gas sensors. *Sens. Actuators B Chem.* 16, 390–395.

Vlachos, D. S., Papadopoulos, C. A., and Avaritsiotis, J. N. (1997). Characterization of the catalyst-semiconductor interaction mechanism in metal-oxide gas sensors. *Sens. Actuators B Chem.* 44, 458–461.

Vosshall, L. B., Wong, A. M., and Axel, R. (2000). An olfactory sensory map in the fly brain. *Cell* 102, 147–159.

Wesson, D. W., Carey, R. M., Verhagen, J. V., and Wachowiak, M. (2008). Rapid encoding and perception of novel odors in the rat. *PLoS Biol.* 6, e82. doi:10.1371/journal.pbio.0060082

Yamazoe, N., and Shimanoe, K. (2008). Theory of power laws for semiconductor gas sensors. *Sens. Actuators B Chem.* 128, 566–573.

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.

Received: 30 September 2011; paper pending published: 16 October 2011; accepted: 02 December 2011; published online: 03 January 2012.

Citation: Yamani JA, Boussaid F, Bermak A and Martinez D (2012) Glomerular latency coding in artificial olfaction. Front. Neuroeng. ***4****:18. doi: 10.3389/fneng.2011.00018*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 20 December 2011
doi: 10.3389/fneng.2011.00016

An investigation on the role of spike latency in an artificial olfactory system

Eugenio Martinelli[1†], Davide Polese[1†], Francesca Dini[1], Roberto Paolesse[2], Daniel Filippini[3], Ingemar Lundström[3] and Corrado Di Natale[1]*

[1] *Department of Electronic Engineering, University of Rome Tor Vergata, Roma, Italy*
[2] *Department of Chemical Science and Technology, University of Rome Tor Vergata, Roma, Italy*
[3] *Division of Applied Physics, Department of Physics, Chemistry and Biology, Linköping University, Linköping, Sweden*

Edited by:
Ramon Huerta, University of California, USA

Reviewed by:
Jordi Fonollosa, Universitat de Barcelona, Spain
Alex Vergara, University of California, USA

****Correspondence:***
Corrado Di Natale, Department of Electronic Engineering, University of Rome Tor Vergata, Via del Politecnico 1, 00133 Roma, Italy.
e-mail: dinatale@uniroma2.it

[†]Eugenio Martinelli and Davide Polese have contributed equally to this work.

Experimental studies have shown that the reactions to external stimuli may appear only few hundreds of milliseconds after the physical interaction of the stimulus with the proper receptor. This behavior suggests that neurons transmit the largest meaningful part of their signal in the first spikes, and than that the spike latency is a good descriptor of the information content in biological neural networks. In this paper this property has been investigated in an artificial sensorial system where a single layer of spiking neurons is trained with the data generated by an artificial olfactory platform based on a large array of chemical sensors. The capability to discriminate between distinct chemicals and mixtures of them was studied with spiking neural networks endowed with and without lateral inhibitions and considering as output feature of the network both the spikes latency and the average firing rate. Results show that the average firing rate of the output spikes sequences shows the best separation among the experienced vapors, however the latency code is able in a shorter time to correctly discriminate all the tested volatile compounds. This behavior is qualitatively similar to those recently found in natural olfaction, and noteworthy it provides practical suggestions to tail the measurement conditions of artificial olfactory systems defining for each specific case a proper measurement time.

Keywords: spike latency, spiking neural networks, artificial olfactory systems, color indicators

INTRODUCTION

The processing of signals from sensorial inputs is an important function in all living beings. In most multicellular animals this process is carried out by the nervous system that is formed by a dense network of specialized cells called neurons. Experimentally gained evidences about the functionalities of the nervous system led to the design of mathematical models with the twofold purpose to elucidate the physiological processes and also to artificially reproduce the natural functions. During the years, these models, broadly called artificial neural networks, have progressively extended the similarity with Nature including functions and structures as those became known by physiological investigations. To this regard, neural networks involving spiking neuron units raise the level of biological similarity (Maass, 1997, 1999), incorporating the spatio-temporal computation (Ferster and Spruston, 1995). Although the behavior of individual neurons is well characterized by phenomenological models (Gerstner and Kistler, 2002), it is not yet completely clear how the information is distributed in the spike patterns and which code neurons use to transmit this information. To this end, behavioral studies show that the reaction times of several animals to external stimuli can be surprisingly short (Rieke et al., 1996). Moreover, recent findings evidences that neurons in the cortex can perform very fast analog computations. For example, humans can analyze and identify visual inputs in less than 100 ms. Then, considering that this process involves at least 10–15 synaptic steps from the retina to the temporal hemisphere, less than 10 ms is the remaining time for the neuron processing (Thorpe et al., 1996). These evidences are not compatible with the analysis performed with standard descriptors like firing rate (Michael and Johnson, 2003) but they are compatible with alternative descriptors, such as spike latencies. This does not mean that the rate coding is not used, instead that when fast processing is needed, latency coding schemes are preferred (Maass, 1999). Spike latency is defined as the time interval between the application of the stimulus and the first spike.

The importance of the processing of this feature is found in more brains zones like the visual or auditory area (Heil, 2004; Gollisch and Meister, 2008). The very fast discrimination and recognition of odors shown by several animals to particular stimuli suggests that this feature can play a significant role also in olfaction (Galizia and Menzel, 2000; Uchida and Mainen, 2003). This conjecture is supported by the recent observation that raise the level of biological similarity latency patterns of olfactory glomeruli contain a quantity of information that is sufficient for higher brain centers to identify odors and their concentrations (Junek et al., 2010).

The features of spike encoding attracts also researchers investigating artificial senses as a step toward the implementation of biological computational paradigms.

The studies on artificial senses are rather advanced for "physical" senses such as sight and hearing, while the artificial analog of

those senses involving chemistry, such as olfaction and taste, are still rather primitive.

Artificial olfaction stemmed at the end of the eighties from the observation that solid-state gas sensors and olfactory receptors (in amphibians, insects, and mammals) are characterized by a wide receptive field (Sicard and Holley, 1984; Malnic et al., 1999; Hallem et al., 2004). This conjecture evolved in the electronic nose concept that is defined as an ensemble of partially specific chemical sensors complemented by some pattern recognition algorithm (Persaud and Dodds, 1982; Gardner and Bartlett, 1994).

Besides the combinatorial sensitivity of receptors olfaction reveals a complex structure that is supposed to play a relevant part in odor recognition (Firestein, 2001). Olfaction is based on few hundreds of different receptor classes expressed by millions of olfactory neurons organized in the olfactory epithelium at the interface with the external environment. The large redundancy of the epithelium, namely many olfactory neurons carrying the same receptors, is used by the convergence of olfactory neurons signals into structures called glomeruli that are organized in the olfactory bulb. Many experimental evidences support the common opinion that each olfactory neuron expresses only one type of receptor and that each glomerulus receives signals from olfactory neurons carrying the same type of receptor (Imai et al., 2006).

The features of spike encoding attracts also researchers investigating artificial olfaction as a step toward the implementation of biological computational paradigms (Martinelli et al., 2006; Chen et al., 2011). In these works integrate and fire models were optimized with the purpose of increasing odor recognition. A more complex approach was made by Koickal et al. (2007) who implemented in a chip a bio-inspired signal processing of the sensor signals and indicated with simulated data the potential impact of spike latency in odor identification. Noteworthy, all these studies were concerned with arrays of few sensor units (where each sensor actually plays the role of an artificial glomerulus). The discrimination content of the latency coding in natural olfaction is then a very interesting aspect that could also produce remarkable advantages in artificial olfaction. Eventually, the application of this concept could indeed allow reducing the measurement time with a consequent minor exposure of the sensors to the samples and probably a reduced rate of poisoning of the sensor surface. On the other hand, the exploitation of latency requires the use of a sensor system incorporating the main features of the natural olfaction. These are the very large number of individual sensors, the glomeruli layer and the spike encoding and processing of the sensor signals. The difficulties to design and assemble large sensor arrays are brilliantly solved by optical sensors. Several years ago it was demonstrated that an image sensor (such as a CCD camera) could conveniently measure, at once, the optical changes occurring in a large number of fluorescent indicators deposited on the tips in a bundle of optical fibers (Dickinson et al., 1999).

The basic property of an image sensor is the segmentation of a whole scene into a number of elementary units, called pixels. Each pixel corresponds to one photo detector measuring the light intensity shining from a section of the whole scene. Eventually, when an image sensor captures a sensitive surface coated by a continuous layer of chemical indicators, the sensing layer is segmented into a number of elementary units corresponding to the pixels of the image. Then, since it is possible to evaluate the optical properties of single pixels, each pixel of the image may correspond to an individual sensor. To this end, even low-resolution images may result in thousands of independent sensing units, then under the hypotheses that different types of indicators are optically different (this trivially means that each indicator has a different color) it is possible to cluster the camera evaluation of the color of individual pixels in abstract classes each containing pixels carrying the same chemical indicator (Di Natale et al., 2008).

In this paper, a spiking neural network, mimicking the first signal elaboration of the glomerular layer spikes sequences, has been studied with data generated by an artificial olfactory platform based on the principle described above.

The platform is formed by an artificial epithelium composed of a layer of chemical indicators. The concept is illustrated in **Figures 1** and **2**. The sensing layer is illuminated by a programmable computer screen and imaged by a digital camera (Filippini et al., 2006). Each single pixel composing the image is then considered as an independent sensor characterized by a proper optical set of features given by the chemical reporters spotted on its area (Di Natale et al., 2008). Then, by exploiting their optical signature, the pixels can be unsupervisedly grouped in classes giving rise to an analog of the Olfactory Receptors Neurons–Glomerulus relationship (Korsching, 2002). This platform can then efficiently mimic the dynamics between signals of single receptors and how these are transformed by the convergence into the glomeruli. Hardware development is limited to the receptor units (the color indicators) while glomeruli are a software implementation. Post-glomeruli processing can be easily accomplished via software and results can be strongly connected to the physical and chemical properties of the interaction of volatile compounds with the receptor layer.

In this way, it is possible to define a processing architecture that describes the pixels as artificial olfactory receptor neurons and the convergence classes as a unit, where the mean signals of the afferent artificial olfactory receptor neurons are the output signals of artificial glomeruli. This architecture that preserves many features of its biological counterpart offers also interesting practical advantages such as the largest signal to noise ratio of glomeruli with respect to the signals of individual sensors and a very large tolerance to single sensors fault events.

Glomeruli output signals are encoded into spike sequences and then utilized as input to a spiking neural network. The latency of the first two spikes of the neurons of the spiking neural network has been considered as the network output, and the application of the principal component analysis (PCA) to these quantities revealed the capability of latencies to discriminate between different volatile compounds.

MATERIALS AND METHODS

A spiking neural network has been applied to process the data produced by the artificial olfactory platform described above. The scheme of the whole system is shown in **Figure 3**. In order to feed the spiking neural network in a way similar to natural systems, the glomerular units besides to average all the signals from their afferent artificial olfactory neurons (corresponding to the individual pixels of the image sensor) have to encode the analog signals into spikes. The spiking network is formed by a single

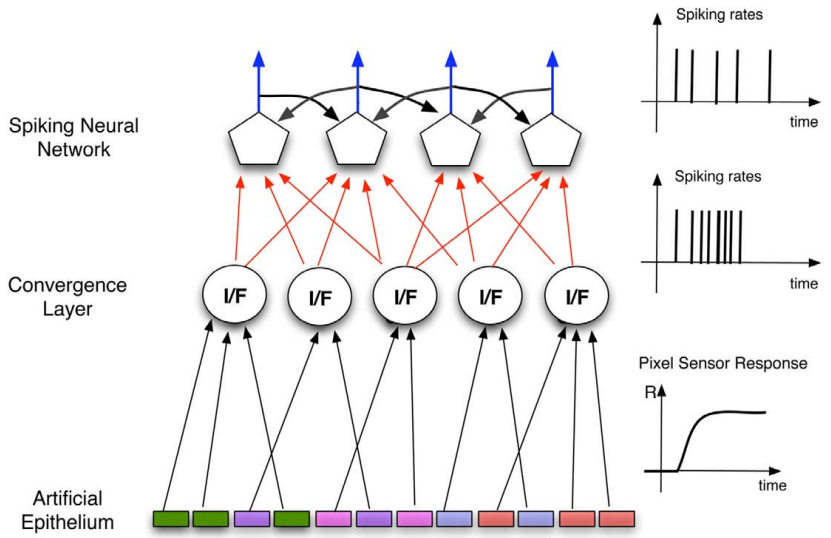

FIGURE 1 | Schematics of the architecture of the Artificial olfactory system. The pixels defining the indicator spots, forming the artificial epithelium, converge, according to their optical signatures to the units of a convergence layer mimicking the role of the olfactory bulb. According to this analogy the units in this layer can be considered as artificial glomeruli. In this layer the average of the signals of the afferent pixels is calculated and then encoded as a spike sequence that is utilized as the input to the processing spiking neural network.

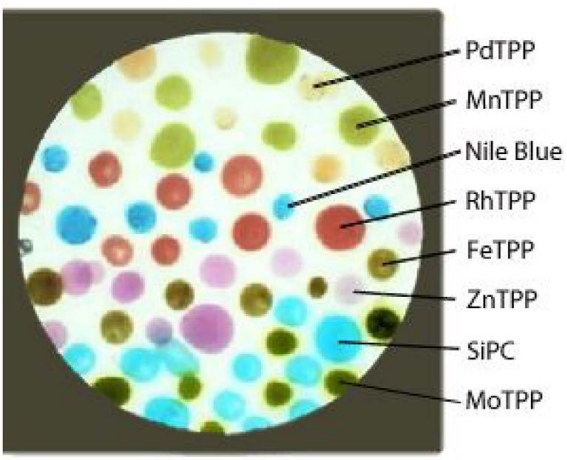

FIGURE 2 | Picture of the artificial epithelium, differently colored spots are the indicators characterized by different hues. The colorless background is coated by the supporting polymer.

layer. Each neuron of the network receives the excitatory inputs from the glomeruli and it is contemporaneously inhibited by the two adjacent neurons.

ARTIFICIAL EPITHELIUM

The sensing layer was formed by eight molecular indicators. Six of them were metal complexes of the (5,10,15,20-tetraphenylporphyrin), namely (5,10,15,20-tetraphenylporphyrin) palladium [PdTPP], (5,10,15,20-tetraphenylporphyrin)rhodium [RhTPP], (5,10,15,20-tetraphenylporphyrin)zinc [ZnTPP], (5,10,15,20-tetraphenylporphyrin)molybdenum [MoTPP], (5,10,15,20-tetraphenylporphyrin)iron [FeTPP], (5,10,15,20-tetraphenylporphyrin)manganese [MnTPP], a Silicon (phthalocyanine) [SiPC], and finally a pH indicator known as Nile Blue (NB). The sensing molecules were dispersed in a membrane based on plasticized poly(vinyl chloride) (PVC; membrane composition in weight: 33% PVC, 66% bisethylhexyl sebacate and 1% indicator). The sensing layer was prepared onto a 25-mm diameter transparent substrate, a Thermanox plastic coverslip provided by Nunc. The polymer membrane was at first casted onto the coverslip to fully coat the surface. Then each indicator-polymer membrane was randomly spotted in several replicas, in a number variable from 6 to 11, above the polymeric coating, to obtain the appearance shown in **Figure 2**. The polymer layer has the main function to support the sensing molecules avoiding the formation of indicator aggregations, maintaining the molecules in a semi-solvent condition.

MEASUREMENT SETUP

Absorbance properties of the sensing spots were measured according to the computer screen photo-assisted technique (CSPT) where a computer screen is used as light source and a digital camera is used as a detector (Filippini et al., 2003). It was shown that such an arrangement, even if based on low-cost components, has enough sensitivity to capture the changes of the optical properties occurring in layers of metalloporphyrins when these are exposed to volatile compounds (Filippini et al., 2006) with a sensitivity comparable with that exhibited by solid-state sensors such as quartz microbalances (Di Natale et al., 2010). Experiments were carried out with a computer screen (Philips 170S4) and a webcam (Philips SPC650NC/97). The camera was operated at a resolution of 160×120 pixels, the signal intensities were given in camera units according to the 8-bit analog to digital conversion of the light intensity embedded in the camera driver. The arrangement

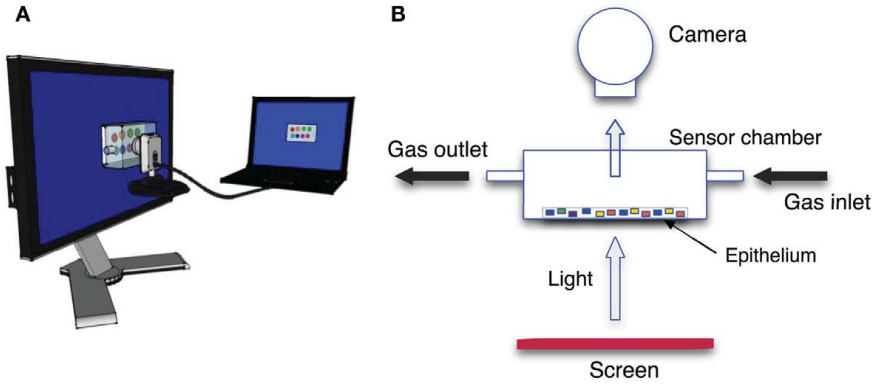

FIGURE 3 | Conceptual scheme of the CSPT setup. (A) Both the screen and the camera are connected to the same computer where all the measurements are controlled and data are recorded. **(B)** The sensing layer is enclosed in a gas-tight and transparent cell where vapors can be passed. The cell is leaned over the surface of a LCD screen.

of the experimental setup is shown in **Figure 3**. The plastic substrate was placed in a sealed cell with transparent windows in order to be probed by light. Organic vapors diluted in a nitrogen flow were passed through the cell. The tested volatile organic compounds (VOCs) were two alcohols (ethanol and butanol) and three amines (trimethylamine, triethylamine, and butylamine). Mixtures of ethanol and butanol, trimethylamine and triethylamine, and trimethylamine and ethanol have also been tested. Besides trimethylamine, all compounds are liquid at standard pressure and temperature, then the gas samples were obtained diluting the saturated pressure at room temperature in a pure nitrogen gas flow. Mass flow controllers regulated the dilution factor and the total flow during the measurements. The same dilution factor (5%) was applied for all the compounds, due to the different phase change parameters, the concentration of the different compounds covered an interval from 1100 ppm for butylamine to 9000 ppm for ethanol. Trimethylamine was measured from a certified bottle where it was diluted in pure nitrogen gas at the concentration of 500 ppm. Each vapor was measured in triplicate. For each vapor the sensing layer was exposed for 400 s and then kept 1100 s under pure nitrogen gas flow to purge the indicators and to restore the initial conditions.

These compounds are moderate Lewis acids and bases that are known to elicit color change in porphyrins and acid-base indicators. The opposite behavior of amines and alcohols provides a significant test where the straightforward discrimination of amines from alcohols is complemented by the recognition of the subtle differences between compounds inside each family. Binary mixtures were also tested to study the interference of compounds belonging to the same family or to different families. In order to evaluate also the reproducibility of the sensor system each vapor and mixture was measured in triplicate.

The reaction to gas was probed with a pure green light illumination. As known from previous investigations, the wavelengths contained in this color are suitable to appreciate the spectral changes occurring in all the indicators. Hence the sensing layer was illuminated with a pure green light, obtained with the RGB code: [0 255 0]. As a consequence, the intensity of the camera green channel of each pixel was the sensors output signal. During the exposure to gas, the camera took a still image each 5 s.

GLOMERULI LAYER DEFINITION AND ANALOG-TO-SPIKING CONVERSION

As previously mentioned, CSPT gives the opportunity to introduce an elegant and simple methodology to reproduce the connection between the olfactory receptors and the glomeruli. Under the hypothesis that the indicators are characterized by different colors, it is possible comparing the color measured in each pixel to assign pixels to classes whose elements are pixels carrying similar indicators. In this way, the physical pixels are the artificial olfactory neurons and the abstract classes are the olfactory glomeruli. CSPT offers a powerful method for color measuring, it is based on the formation of a fingerprint obtained exposing the sample to a set of colors achieved blending the RGB values regulating the LCD screen color, and measuring the resulting appearance of the sample in the three channels (red, green, and blue) of the digital camera. This method was demonstrated to be able to discriminate between subtle color changes in colorimetric tests (Filippini and Lundström, 2006). For the scope of glomeruli definition the sensing layer was illuminated by the computer screen programmed to display a rainbow of 50 colors from purple to red. The layer was imaged by the webcam encoding the color in an 8-bit scale separated in the red, green, and blue channels. The signals from the red, green, and blue channels were concatenated, in this order, to form a fingerprint vector. As a result, the color of pixels was encoded in a fingerprint vector composed by 50 illumination colors $\times$ 3 camera color channels. **Figure 4** shows the collected fingerprints. In the ideal case of a perfectly homogeneous distribution of the indicators and a uniform screen illumination, only nine different fingerprints (corresponding to the eight indicators and the polymer substrate) should be visible in **Figure 4**. Actually, the fingerprints are almost continuously distributed evidencing a high level of non-homogeneity in sensing spot formation. The clustering of the fingerprints of **Figure 4** is in principle a straightforward supervised operation. Indeed, the position of each spot is known and indicators can be easily identified by their color. However,

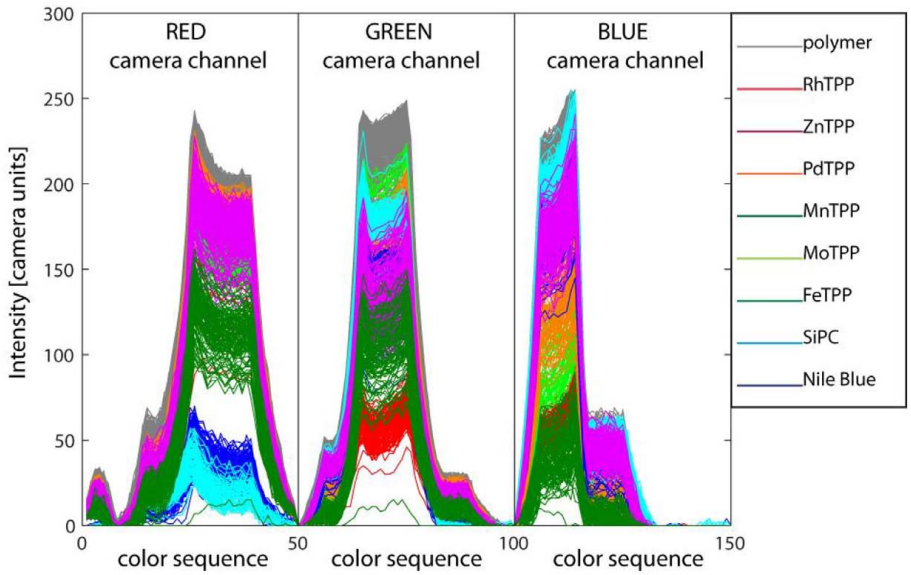

FIGURE 4 | Computer screen photo-assisted technique fingerprints of all the pixels imaged in Figure 2. The fingerprints are formed concatenating the red, green, and blue channels levels read by the camera under the exposure to a rainbow of 50 colors from purple to pure red. The fingerprints are differently colored according to the indicator imaged in the corresponding pixel. Straightforwardly, the largest signals are obtained for the pixels were the pure supporting polymer is imaged.

it is interesting to incorporate in the artificial platform architecture an automatic procedure clustering together pixels imaging the same kind of indicators. This gives rise to an elegant mimic of natural olfactory neurons where pixels are the olfactory neurons, and the indicators are the olfactory receptor. The receptor itself provides both the chemical sensitivity and the criteria to the convergence to the relevant glomeruli. The efficiency of CSPT to capture colors allows using automatic clustering with a negligible amount of misclassifications. From a biomimetic point of view, clustering of similar pixels provides the same class membership disregarding the arrangement of spots, i.e., the distribution of a given type of olfactory receptor neurons in the epithelium. From a practical point of view this allows for instance a prompt replacement of the artificial epithelium and even a simple communication of signals between different sensing units (Polese et al., 2011).

The fingerprints of **Figure 4** can be conveniently clustered into an arbitrary number of classes by means of any unsupervised classification algorithm. Here a simple K-nearest neighbor (KNN) algorithm has been chosen (Duda et al., 2000). According to the practical implementation of the artificial epithelium, shown in **Figure 2**, the indicators cover only a portion of the sensing layer. The rest of the area is coated with the supporting polymer and its color is clearly different from the color of the indicators, this is quite visible in **Figure 2** where the polymer-coated pixels are clearly separated from the dyed pixels. In order to limit the analysis to the pixels carrying information about the indicators, the polymer-coated pixels were segregated applying a two classes KNN classifier. Here the obvious choice of two classes was made possible by the straightforward difference between the fingerprints of dyed and non-dyed pixels. However, the number of classes is a free parameter of unsupervised classifiers, and the number of classes has to be chosen according to the specific conditions of each application. In this case, the number of classes can be chosen in order to abide the biological paradigm according to which in a single glomerulus only olfactory neurons expressing the same receptor converge. This means that in the artificial system the classes defined by KNN have to contain only pixels carrying the same (or very similar) indicators. In order to fulfill this requirement, a number of classes greater than the actual number of indicators is necessary and here, 13 classes have been considered. This number, empirically chosen, is large enough to avoid misclassifications, and at the same time, it is sufficiently small to guarantee that each class collects a sufficient number of artificial olfactory neurons adequate to provide, through the signal averaging, a significant increase of the signal to noise ratio. Besides averaging the incoming pixel signals, glomeruli are requested to encode the signal into spikes in order to be processed by the spiking neural network. Here, an "integrate and fire" (I/F) algorithm was used (Gerstner and Kistler, 2002). The algorithm is conveniently described by the equivalent electronic circuit shown in **Figure 5**. The inter-spike time is determined by the magnitude of the analog input signal and the circuit time constant, that is given by the product of the two algorithm parameters: R and C, according to the following equation:

$$u(t) = R \cdot I(t) - C \cdot R \cdot \frac{du(t)}{dt} \tag{1}$$

$I(t)$ is the input signal and $u(t)$, the voltage across the capacitance. The circuit generates a spike anytime the voltage $u(t)$ reaches the threshold value (Vthreshold in **Figure 5**). The critical

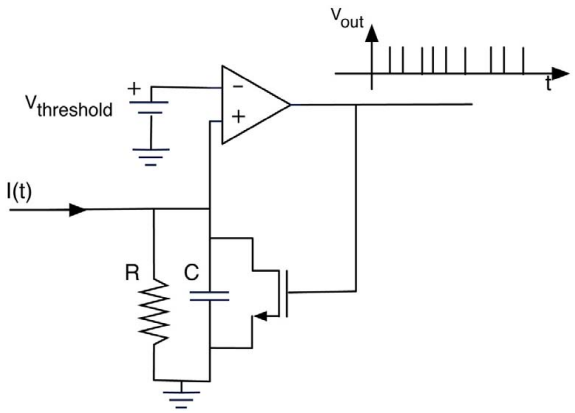

FIGURE 5 | The I/F algorithm implemented in the glomeruli units is represented through its electronic equivalent circuit. The average signal of the pixels pertaining to the glomerulus is represented by the input current $i(t)$ and the output voltage is the encoded spiking sequence. Circuit elements, including the operational amplifier, are ideal and their values are internal parameters of the algorithm.

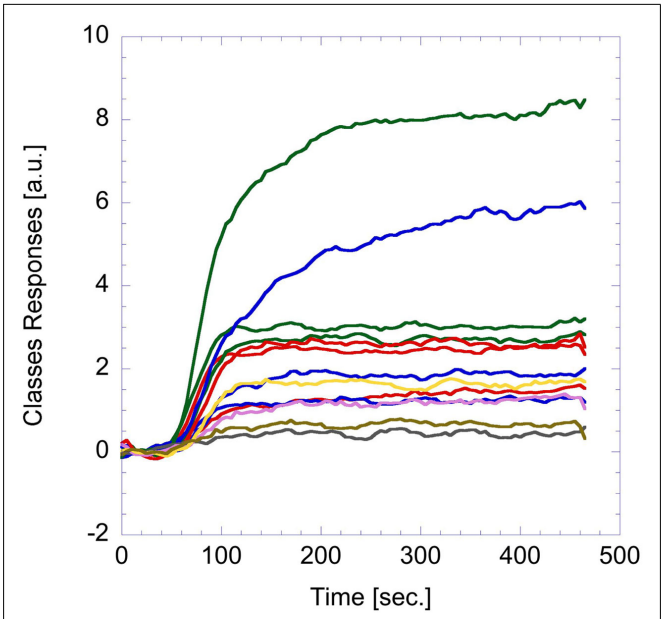

FIGURE 6 | Example of analog glomeruli signals. These signals were caused by the exposure to trimethylamine vapors (about 500 ppm in nitrogen).

parameter of the algorithm is the time constant that is defined as the product of R and C in the analog equivalent circuit. This value has been empirically fixed to 21.3 s to obtain a sufficiently large number of spikes from the glomeruli signals. The same parameters were applied for all the glomeruli. **Figure 6** shows a typical analog glomeruli signal emerging as a consequence of the exposure of the artificial epithelium to a vapor. To obtain only to the contribution of the effect of the gas, the glomeruli signals value recorded immediately before the exposure to gas was subtracted from the recorded signal. The difference signals were then converted into spike sequences and processed by the spiking neural network, according to the procedure outlined in **Figure 1**.

It is important to note that the signals can be either positive or negative indicating that the interaction of the indicator with the volatile compound elicits different changes in the optical spectrum of the indicator including a change of magnitude and a shift of the optical features. A combination of these changes with the spectral response of the camera green channel and the green light spectra of the LCD screen may then result in positive or negative changes of the camera signal. In a few cases a negative change of the intensities in the green camera channel was observed and only for one of the glomeruli. According to Eq. 1, a negative analog signal does not cause a spike train. The glomerulus exhibiting negative signals was therefore not a part of the further processing in the spiking network. Actually also other sensors can present signals that can be either positive or negative, this is the case, for instance, of the metal-oxide semiconductor gas sensors used in a previous study (Chen et al., 2011).

THE SPIKING NEURAL NETWORK

The spiking neural network used to process the artificial olfaction data was composed by neurons arranged in a single layer as shown in **Figure 1**. The neurons were defined according to a phenomenological neuron model (Gerstner and Kistler, 2002) and they were endowed of both excitatory (from the glomeruli) and inhibitory (from the other network neurons) inputs. As a consequence the relative position of the neurons in the network architecture and the inputs distribution into the various neuron units are important. Since, the glomeruli have been defined above as structure-less abstract entities it is necessary to introduce an ordering criteria. As illustrated in Section "Glomeruli layer definition and analog-to-spiking conversion" the color of the indicators is an optimal quantity to identify the indicators in the array, and the glomeruli have been simply defined applying the KNN algorithm to the CSPT fingerprints. Here, the same quantity is further used to define the interface between the glomeruli units and the spiking neural network. Each glomerulus is defined by its KNN centroid vector, roughly corresponding to the average of the CSPT fingerprints. The connection between glomeruli and neural network neurons has been based on a hierarchical clustering of the centroid fingerprint vectors. The outcome of the process is illustrated in **Figure 7A** where the 13 centroid fingerprints are shown in the plane of the first two principal components of the PCA calculated on the centroid fingerprints matrix. In figure the mapping between the glomeruli and the neural network neurons is also visible. All the glomeruli signals contribute to their relevant neuron with an excitatory input. The number of network neurons has been fixed to 12. With this number of neurons the excitatory inputs are between 2 and 4 for each neuron. In **Figure 7B** the relationship between the chemical indicators, the glomeruli, and the neural network neurons is shown. In this figure the neurons are ordered according to a hierarchical clustering criteria visible in **Figure 7A**. This order is important because each neuron receives inhibitory inputs from the first neighbors in

the list shown in **Figure 7B**. Finally it is worth the remark that since the chemical sensitivity of the indicator and its color are not correlated, no chemotopic feature is expected to be found in the neural network structure. According to the phenomenological neuron model, the neuron state is controlled by the action potential (AP):

$$AP_j(t) = \sum_{i \in \text{input}} w_{ij} \cdot \text{kernel}(t - t_n) \cdot H(t - t_n) + \sum_{t_m} \text{Refractory}(t - t_n) \cdot H(t - t_n) \quad (2)$$

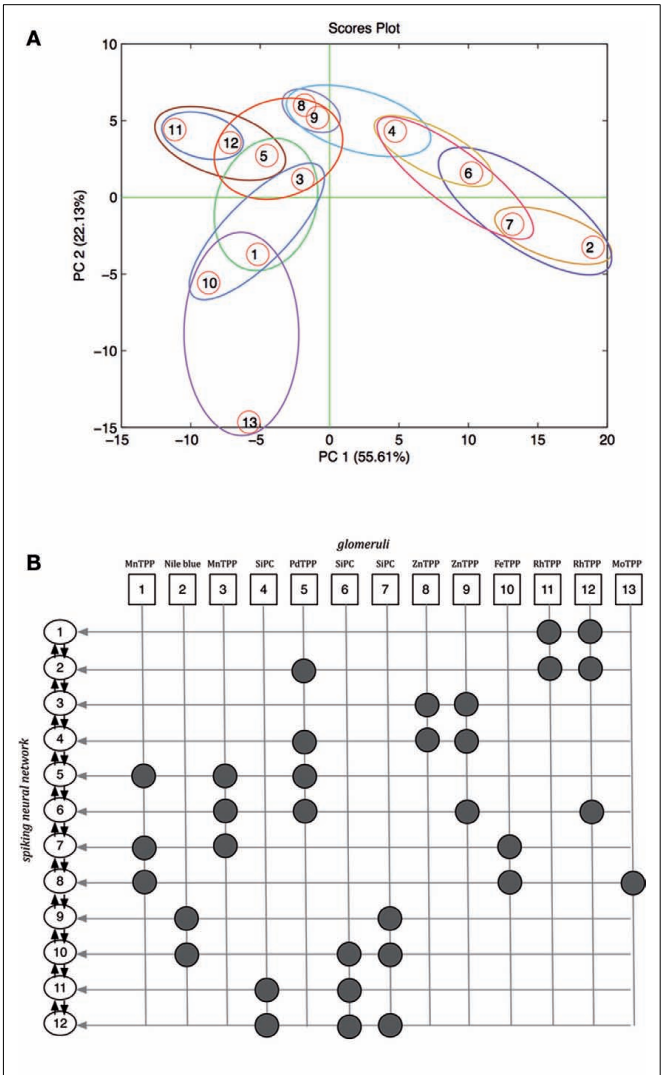

FIGURE 7 | (A) Scores plot of the first two principal components of the PCA of the matrix of the centroid fingerprints as defined by the application of the KNN to the fingerprints matrix. Closed regions show the membership of glomeruli to the network neurons. **(B)** Relationship between chemical indicators, glomeruli, and neural network neurons excitatory inputs. Glomeruli signals contribute as excitatory inputs to the spiking neural network neurons. The arrows between the network neurons indicate the inhibitory signals.

where $H(t–t_m)$ is the Heaviside function. When a pre-synaptic spike occurs at the time $t = t_n$, it produces an AP variation proportional to the kernel function through the pre-synaptic weight (w_{ij}). Once the AP reaches the threshold (θ) at the time $t = t_m$ the neuron fires and the AP is reset to zero by the refractory function. The expression of the learning, the kernel and the refractory functions are shown in **Table 1**. The excitatory weights are trained with a Hebbian learning rule where the change, δw in the upgrade coefficient, w, observes the following equation:

$$\delta w = \eta \cdot \left[\sum_{t_i} w_{in} + \sum_{t_0} w_{out} + \sum_{t_i, t_0} w(t_0 - t_i) \right] \quad (3)$$

where w_{in} is the discrete increment of the input weight corresponding to one input spike, w_{out} is the discrete decrement of the input weight corresponding to one output spike, and $w(t_{0–ti})$ is the learning function, t_i and t_0 are the input and output firing time and η is the learning rate.

The initial values of the synaptic weights of the excitatory input are randomly generated in the range (0, 1) and the APs were initially set to zero. For the inhibitory connections, two configurations of synaptic coefficients have been considered in the analysis. In the first configuration, the coefficients have been randomly generated in the range (−3, −0.5) and in the second case the contribution of the inhibition has been removed setting all the inhibitory synaptic weights to zero. In both cases these coefficients have been maintained constant during the following network training. The neural network was unsupervisedly trained presenting 100 times, in a randomized order, the spiking sequences encoding the artificial glomerular responses to the tested vapors. In order to evaluate the performance of the complete artificial olfactory system (composed by the physical layer of receptors, the abstract glomeruli layer, and the spiking neural network) to discriminate the tested volatile compounds, the output of the network have been described by two features describing the short time response and the long time response of to the applied stimulus, respectively. The short time response is considered by the latencies of the first two spikes of each neuron, while the long time response is explained by the neurons firing rate averaged in the time window corresponding to the application of the stimulus. To study the contribution of the spiking neural network

Table 1 | List of the functions describing the phenomenological neuron model used in the spiking neural network.

Learning function	if $\Delta t \leq 0$	$W(\Delta t) = \eta \cdot \exp(\Delta t / t_{syn}) \times [A_+(1 - \Delta t/\tilde{\tau}_+) + A_- \cdot (1 - \Delta t/\tilde{\tau}_-)]$
	$f\ \Delta t > 0$	$W(\Delta t) = \eta \cdot \left[-A_+ \cdot \Delta t/\tilde{\tau}_+ - A_- \cdot \Delta t/\tilde{\tau}_-\right]$
Kernel function		$\frac{1}{1-\frac{\tau_m}{\tau_s}} \cdot [exp(-t_i/\tau_s) - exp(-t_i/\tau_m)]$
Refractory function		$-\vartheta \exp(-t_r/\tau)$

The different τ's appearing in the definitions are time constants, $A_\pm$ are constant parameters and θ is the neuron threshold.

to the overall discrimination capabilities of the whole olfactory system, two additional features describing the glomeruli signals have also been considered. Also in this case a short time and a long time response have been examined. The first feature was calculated as the analog glomeruli signals at the time when the slowest neuron of the spiking network fires its second spike, while the long time response is simply the maximum signal showed by glomeruli during the stimuli application. **Figure 8** illustrates the definition of these features. In correspondence to the exposure to each volatile compound and for each feature, a vector composed by the ordered sequence of features was arranged. The whole set of experiment resulted then in matrices whose elements were the above illustrated features. The matrices were then analyzed with (PCA; Joliiffee, 2002). The main advantage of PCA, in this case, is the possibility to plot in the principal components plane the multidimensional datasets, in this way a simple and effective evaluation of the discrimination capability of each feature can immediately be obtained by a simple visual inspection of the plots.

RESULTS AND DISCUSSION

It is important to remark that this paper does not intend to provide a faithful replica of the olfactory circuits but rather it is centered on the study of the information content of the spiking sequence and in particular of the spikes latency coding. Then the complexity of the network was kept as low as possible adopting an architecture with a single layer of spiking neurons. Two different topologies characterized by the presence and the absence of lateral inhibitions were taken into consideration. The object of the paper was the investigation of the vapor recognition properties of the spike latency. The latency time of the first spike indicated a limited gas discrimination while a significant improvement was obtained considering as the measurement descriptor the latency times of the first two spikes fired by each neuron of the network (i.e., T1 and T2 in **Figure 8**). All glomeruli generates their first two spikes in less than 120 s providing about 65% of reduction of measuring time with respect to the standard measurement protocol where the end of measure occurs when the analog signals reach their steady-state values.

The ensemble of the latency times were joined to form a vector and the matrix collecting the whole experiment was processed by PCA.

Figure 9 shows the plot of the first two principal components of the latency feature of the spiking neural network trained with the lateral inhibition. All the VOCs are discriminated and only a little overlap between Triethylamine and Triethylamine–Trimethylamine mixture is present. This is expected because of the chemical similarity between these compounds. The presence of the lateral inhibition increasing the differences among the neurons distributes the information carried by the glomeruli into the network neurons. As a partial proof, it was observed that without inhibition (**Figure 10**) the discrimination performance

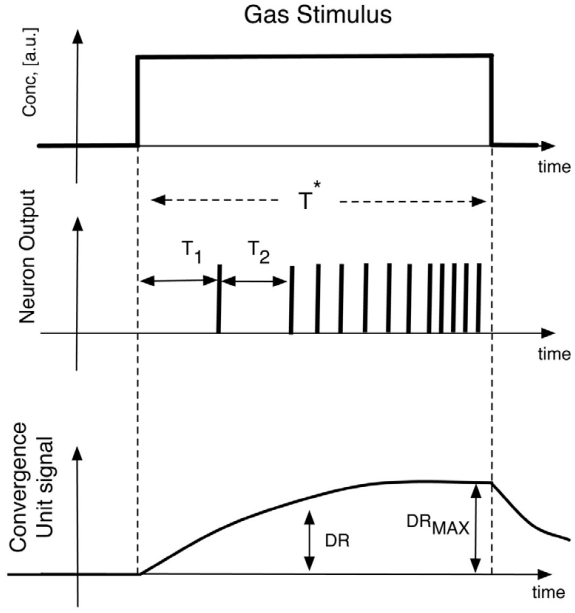

FIGURE 8 | Illustration of the features used to describe both the spiking neural network and the glomeruli layer response to vapors. The time at which DR is calculated corresponds to the time when the slowest neuron fires its second spike. The average firing rate is calculated in the time interval T*.

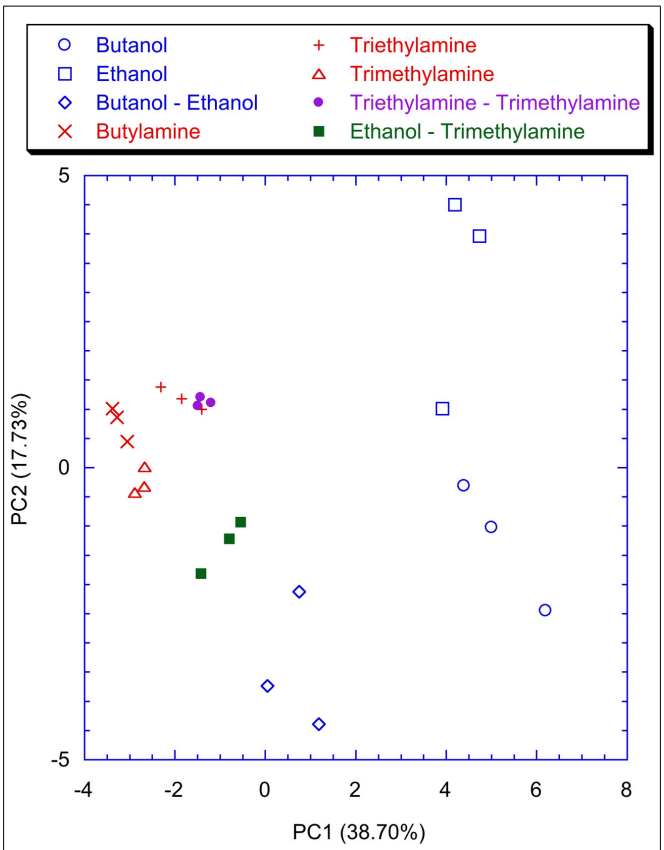

FIGURE 9 | Scores plot of the first two principal components of the PCA of the matrix of latencies of the first two spikes of each neuron of the network trained with the lateral inhibition. Labels indicate the measured vapors.

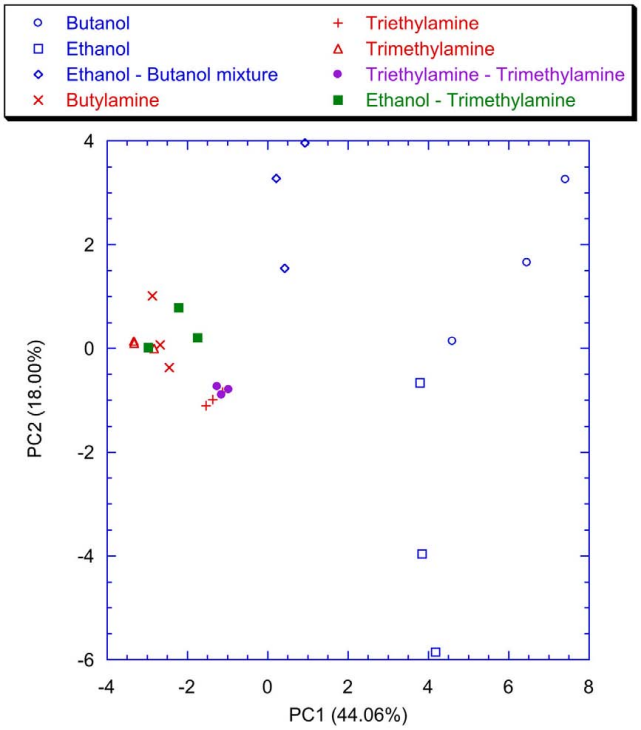

FIGURE 10 | Scores plot of the first two principal components of the PCA of the matrix of latencies of the first two spikes of the neurons of the network trained without the lateral inhibition. Labels indicate the measured vapors.

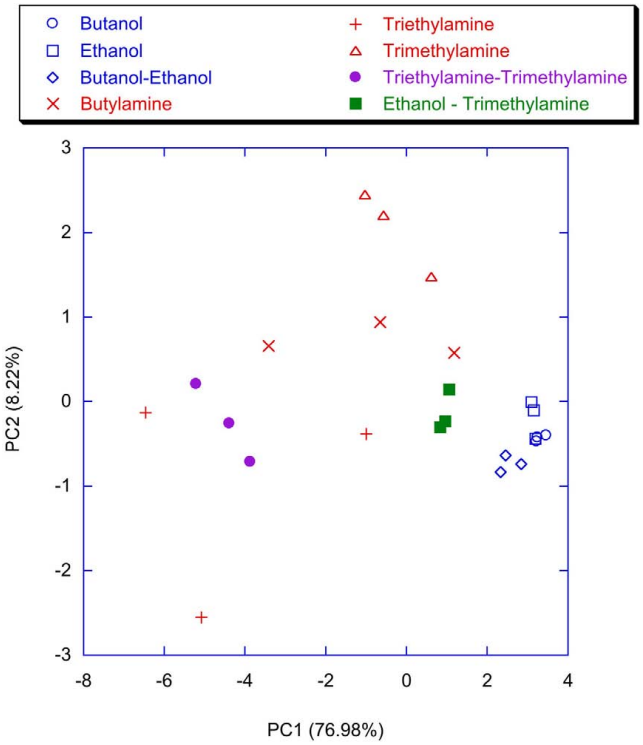

FIGURE 11 | Scores plot of the first two principal components of the PCA of the matrix of the analog signals of the glomeruli evaluated at the time when the slowest neuron fires the second spike. The network was trained with the lateral inhibition and for each measure presentation, the coefficient updating procedure was stopped when the slowest neuron fires the second spike. Labels indicate the measured vapors.

was reduced to the basic, and somewhat chemically straightforward, separation between alcohols and amines. Furthermore, if a certain separation among alcohols can actually be observed the amines are completely overlapped in **Figure 10**. This result is somewhat unexpected because the response of the individual receptors to amines is much larger with than to alcohols. This suggests that the absence of inhibition makes the network unable to capture the differences among different gases when a large response is obtained for the chemical indicators. It is interesting to note that a larger analog signal corresponds to an increase of the spike rate from the convergence layer. The saturation effect could only be related to the frequency response of the network and the I/F algorithm. It could probably be avoided by a careful choice of the parameters. Nonetheless it is interesting to observe that inhibition prevents the occurrence of the saturation effects maintaining unaltered the olfactory system properties in the whole range of signals even for a non-optimized choice of parameters.

To point out the contribution of the spiking neural network, the classification properties of the analog glomeruli signals was considered. In **Figure 11** the plot of the first two principal components of the PCA of the matrix of the glomeruli analog signals is shown. To consider the short time response, the glomeruli signal was considered at the time when the slowest neuron of the spiking neural network, trained with lateral inhibition, fires its second spike. According to **Figure 11** the discrimination of volatile compounds appears worse with respect to that obtained with the spike latencies of the network with inhibition. Another interesting detail to note is that the explained variance of the first two principal components of the PCA of the latencies matrices (**Figures 9** and **10**) are significantly less with respect to the first two principal components of the PCA of the matrix collecting the analog glomeruli signals (**Figure 11**). This result indicates that the spiking network performs a faster separation of the different sources of information; furthermore, in presence of lateral inhibition the differences among neurons responses are more prominent and then the signal decorrelation becomes more evident. However, in the experiment described here the contribution to the classification of the principal components of order higher than two is negligible. It is also important to note that for the network with inhibition, the odor recognition property does not change significantly considering further latencies related to spikes beyond the second spike (data not shown). The classification improves but not dramatically when the average firing rate is considered (see **Figure 12**). These results are qualitatively similar to that obtained using the maximum glomerular signal (see **Figure 13**) suggesting that when the sensors reach the dynamic equilibrium with the gas, the classification is completely explained in the input data and the contribution of the network processing becomes negligible. It is interesting to note that for the average firing rate the

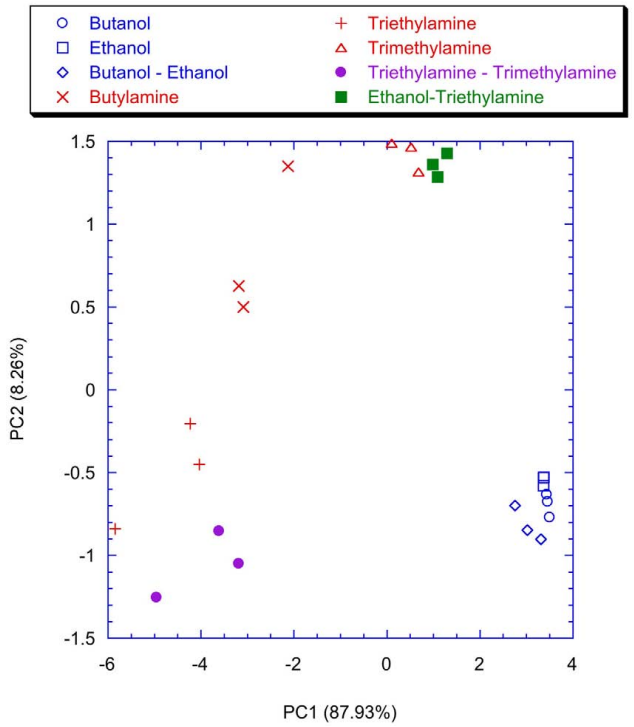

FIGURE 12 | Scores plot of the first two principal components of the PCA of the matrix of the average firing rate of the neurons of the network trained with the lateral inhibition. Labels indicate the measured vapors.

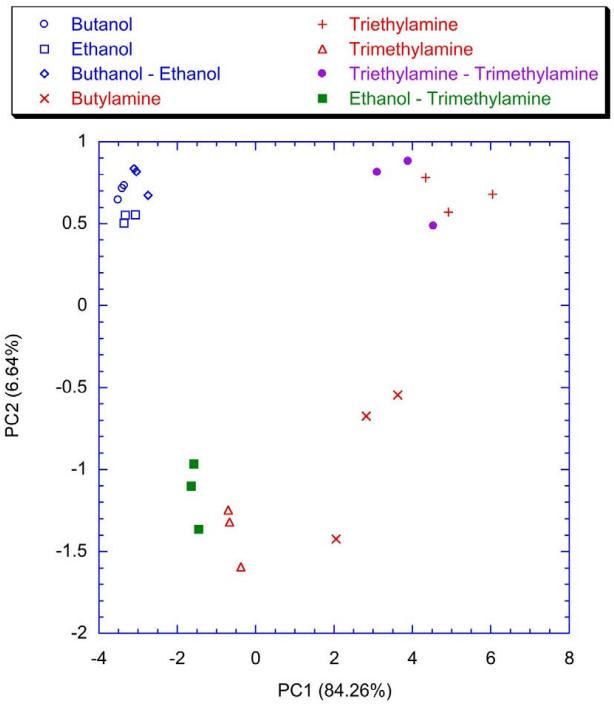

FIGURE 13 | Scores plot of the first two principal components of the PCA calculated on the matrix of the maximum analog signals of the glomeruli. Labels indicate the measured vapors.

variance explained by the first two principal components is more than 95% confirming the large correlation of the neurons firing rates. This result also suggests that the lateral inhibition mainly influences the distribution of the information content carried by the neuron in the initial part of the signal while in the last part its contribution is related to a sort of a scale factor of the glomeruli signals. As further confirmation of this hypothesis, no decrease of performance is observed removing the inhibition using the firing rate as the input (data not shown). This result leads to believe that the lateral inhibition, in this network structure, affects significantly only the information contained in the initial part of the spike sequences. In order to show the importance of the initial part of the response, the network was also trained stopping the learning procedure immediately after the second spike of the slowest neuron is fired. Surprisingly, the results of the PCA calculated with the latencies of the first two spikes and the firing rate are qualitatively similar. This result indicates that the very initial part of the response is of outmost importance and then only this portion of the response could be used in reducing dramatically the computational and the experimental time. Noteworthy, the length of the initial part is dynamically determined by the network as the time necessary to get two spikes from all the neurons and then this time can be variable according to the magnitude and the dynamics of the receptor responses. Ultimately, for a given set of receptors this time depends on the kind of vapor. The reduction of measurement time using only the first portion of sensor response was attempted in the past as by using the dynamic properties of sensors (Davide et al., 1995; Di Natale et al., 1995; Eklöv et al., 1997; Nakamoto et al., 1997; Gutierrez-Osuna et al., 1999; Muezzinoglu et al., 2009). In these studies the measurement time was generally determined according to a previous knowledge about the dynamics of the sensors. In the present case the duration of the measurement is determined by the network processing the global set of receptors hierarchically arranged in glomeruli-like structures and the length of time necessary for a single measurement is variable for each sample but always less than 35% of the time necessary for the analog signal to reach the steady-state. It is important to remark that the dynamic responses of the chemical reporters in the sensing layer are also affected by their spatial arrangement with respect to the inlet and the outlet of the cell where the sensor layer is accommodated. As a consequence, also the latency pattern of the network is dependent by the position of the chemical indicators giving a further degree of freedom for the system design and optimization. Furthermore, considering that the indicators are immersed in a supporting polymer layer, the latency pattern depends also on the diffusion of volatile compounds through the polymer. This feature was demonstrated to lead to a sort of artificial olfactory mucosa (Dini et al., 2009) mimicking the separation of volatile compounds characterized by different alkyl chain lengths and steric effects (Kent et al., 1996). The possibility to use the latency as features for discrimination tasks can open a different approach to chemical sensing. Actually exploiting these descriptors it would be possible to define the gas exposure as the time necessary to obtain the

occurrence of the first spikes at the network output. This strategy would offer some interesting aspects. The first is related to the measurement time that it should be not fixed but it should adapt itself on the gas under measure. Moreover this approach should guarantee a short measurement time with respect to the standard protocol counteracting the aging and poisoning of the devices that is probably the main cause of the chemical sensor drift (Martinelli et al., 2011). It is important to remark, that these results have been derived from an experiment aimed at measuring a limited set of volatile compounds. However, since the test compounds (amines and alcohols) are good representative of donor and acceptor molecules, they elicit a sufficiently range of responses. The findings here outlined may substantially be corroborated by more extensive measurements with a larger number of repetitions to study the relationship between the network property and the sensors reproducibility.

CONCLUSION

An artificial olfactory system based on a large array of optochemical sensors coupled with a bio-inspired signal processing architecture is shown. The processing strategy took into consideration a convergence layer that mimics the role of glomeruli and encodes the signals of the sensors into sequences of spikes to be processed by a spiking neural network. The aim of this work was to investigate the possibility to use the spike latency of the network outputs as a useful descriptor for the odor recognition. The results show that the first two spikes of neural network neurons contain enough information to discriminate the different vapors confirming the experimental evidences obtained on animals. It has also been put in evidence the fundamental role of the lateral inhibition in the information processing of the very initial part of sensor signals and how it is possible to exploit the characteristics of the latency coding to define an adaptive gas exposure strategy.

REFERENCES

Chen, H. T., Ng, K. T., Bermak, A., Law, M. K., and Martinez, D. (2011). Spike latency coding in biologically inspired microelectronic nose. *IEEE Trans. Biomed. Circuits Syst.* 5, 160–168.

Davide, F., Di Natale, C., D'Amico, A., Hierlemann, A., Mitrovics, J., Schweizer, M., Weimar, U., Göpel, W., Marco, S., and Pardo, A. (1995). Dynamic calibration of QMB polymer-coated sensors by Wiener kernel estimation. *Sens. Actuators B Chem.*27, 275–285.

Di Natale, C., Marco, S., Davide, F., and D'Amico, A. (1995). Sensor array calibration time reduction by dynamic modeling. *Sens. Actuators B Chem.* 24, 578–583.

Di Natale, C., Martinelli, E., Paolesse, R., D'Amico, A., Filippini, D., and Lundström, I. (2008). An experimental biomimetic platform for artificial olfaction. *PLoS ONE* 3, e3139. doi:10.1371/journal.pone.0003139

Di Natale, C., Santonico, M., Paolesse, R., Filippini, D., D'Amico, A., and Lundstrom, I. (2010). Evaluation of the performance of sensors based on optical imaging of a chemically sensitive layer. *Anal. Bioanal. Chem.* 397, 613–621.

Dickinson, T., Michael, K., Kauer, J., and Walt, D. (1999). Convergent self encoded bead sensor arrays in the design of an artificial nose. *Anal. Chem.* 71, 2192–2198.

Dini, F., Martinelli, E., Pomarico, G., Paolesse, R., Monti, D., Filippini, D., D'Amico, A., Lundstrom, I., and Di Natale, C. (2009). Chemical sensitivity of self-assembled porphyrin nano-aggregates. *Nanotechnology* 20, 055502.

Duda, R., Hart, P., and Stork, D. (2000). *Pattern Classification.* New York: John Wiley and Sons.

Eklöv, T., Mårtensson, P., and Lundström, I. (1997). Enhanced selectivity of MOSFET gas sensors by systematical analysis of transient parameters. *Anal. Chim. Acta* 353, 291–300.

Ferster, D., and Spruston, N. (1995). Cracking the neuronal code. *Science* 270, 756–757.

Filippini, D., Alimelli, A., Di Natale, C., Paolesse, R., D'Amico, A., and Lundstrom, I. (2006). Chemical sensing with familiar devices. *Angew. Chem. Int. Ed. Engl.* 45, 3800–3803.

Filippini, D., and Lundström, I. (2006). Preferential color substances and optimized illuminations for computer screen photo-assisted classification. *Anal. Chim. Acta* 557, 240–241.

Filippini, D., Svensson, X., and Lundstrom, I. (2003). Computer screen as a programmable light source for visible absorption characterization of (bio)chemical assays. *Chem. Commun.* 9, 240–241.

Firestein, S. (2001). How the olfactory system makes sense of scents. *Nature* 413, 211–218.

Galizia, G., and Menzel, R. (2000). Odour perception in honeybees: coding information in glomerular patterns. *Curr. Opin. Neurobiol.* 10, 504–510.

Gardner, J., and Bartlett, P. (1994). A brief history of electronic noses. *Sens. Actuators B Chem.* 18, 211–220.

Gerstner, W., and Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity.* Cambridge: Cambridge University Press.

Gollisch, T., and Meister, M. (2008). Rapid neural coding in the retina with relative spike latencies. *Science* 319, 1108–1111.

Gutierrez-Osuna, R., Nagle, H. T., and Schiffman, S. S. (1999). Transient response analysis of an electronic nose using multi-exponential models. *Sens. Actuators B Chem.* 61, 170–182.

Hallem, E., Ho, M., and Carlson, J. (2004). The molecular basis of odor coding in the *Drosophila* antenna. *Cell* 117, 965–979.

Heil, P. (2004). First-spike latency of auditory neurons revisited. *Curr. Opin. Neurobiol.* 14, 461–467.

Imai, T., Suzuki, M., and Sakano, H. (2006). Odorant receptor derived cAMP dignals direct axonal targeting. *Science* 314, 657–661.

Joliiffee, I. (2002). *Principal Component Analysis*, 2nd Edn. Berlin: Springer.

Junek, S., Kludt, E., Wolf, F., and Schild, D. (2010). Olfactory coding with patterns of response latencies. *Neuron* 67, 872–884.

Kent, P. F., Mozell, M. M., Murphy, S. J., and Hornung, D. E. (1996). *J. Neurosci.* 16, 345–353.

Koickal, T. J., Hamilton, A., Tan, S. L., Covington, J. A., Gardner, J. W., Pearce, T. C. (2007). Analog VLSI circuit implementation of an adaptive neuromorphic olfaction chip. *IEEE Trans. Circuits Syst.* 54, 60–73.

Korsching, S. (2002). Olfactory maps and odor images. *Curr. Opin. Neurobiol.* 12, 387–392.

Maass, W. (1997). *The Third Generation of Neural Network Models.* Graz: Technische Universität.

Maass, W. (1999). "Computing with spiking neurons," in *Pulsed Neural Networks*, eds W. Maass and C. M. Bishop (Cambridge: MIT press), 55–85.

Malnic, B., Hirono, J., Sato, T., and Buck, L. (1999). Combinatorial receptor codes for odors. *Cell* 96, 713–723.

Martinelli, E., D'Amico, A., and Di Natale, C. (2006). Spike Encoding of artificial olfactory sensor signals. *Sens. Actuators B Chem.* 119, 234–238.

Martinelli, E., Santonico, M., Pennazza, G., Paolesse, R., D'Amico, A., and Di Natale, C. (2011). Short time gas delivery pattern improves long-term sensor reproducibility. *Sens. Actuators B Chem.* 156, 753–759.

Michael, L., and Johnson, B. (2003). Olfactory coding in the mammalian olfactory bulb. *Brain Res. Rev.* 42, 23–32.

Muezzinoglu, M., Vergara, A., Huerta, R., Rulkov, N., Rabinovich, M., Selverston, A., and Abarbanel, H. (2009). Acceleration of chemosensory information processing using transient features. *Sens. Actuators B Chem.* 134, 507–512.

Nakamoto, T., Okazaki, N., and Morizumi, T. (1997). High speed active gas/odor sensing system using adaptive control theory. *Sens. Actuators B Chem.* 41, 183–188.

Persaud, K., and Dodds, G. (1982). Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature* 299, 352–355.

Polese, D., Magna, G., Dini, F., Martinelli, E., Paolesse, R., D'Amico, A., Filippini, D., Lundström, I., and Di Natale, C. (2011). Facile sensors replacement in optical gas sensors array, Proc. of Eurosensors XXV, Athens (4–6/9/2011) to appear in Procedia Engineering.

Rieke, F., Warland, D., de Ruyter, R., van Steveninck, R., and Bialek, W. (1996). *Spike-Exploring the Neural Code*, Cambridge: MIT press.

Sicard, G., and Holley, A. (1984). Receptor cell responses to odorants: similarities and differences among odorants. *Brain Res.* 292, 283–296.

Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. *Nature* 381, 520–522.

Uchida, N., and Mainen, Z. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat. Neurosci.* 6, 1224–1229.

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.

Received: 19 September 2011; paper pending published: 28 October 2011; accepted: 28 November 2011; published online: 20 December 2011.

Citation: Martinelli E, Polese D, Dini F, Paolesse R, Filippini D, Lundström I and Di Natale C (2011) An investigation on the role of spike latency in an artificial olfactory system. Front. Neuroeng. ***4****:16. doi: 10.3389/fneng.2011.00016*

frontiers in
NEUROENGINEERING

REVIEW ARTICLE
published: 18 April 2012
doi: 10.3389/fneng.2012.00007

Synaptic inhibition controls transient oscillatory synchronization in a model of the insect olfactory system

Collins Assisi and Maxim Bazhenov*

Department of Cell Biology and Neuroscience, University of California, Riverside, CA, USA

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Hari S. Sharma, Uppsala University, Sweden
Thomas A. Cleland, Cornell University, USA
Stijn Cassenaer, California Institute of Technology, USA

****Correspondence:***
Maxim Bazhenov, Department of Cell Biology and Neuroscience, University of California, Riverside, CA 92521, USA.
e-mail: maksim.bazhenov@ucr.edu

In a variety of neuronal systems it has been hypothesized that inhibitory interneurons corral principal neurons into synchronously firing groups that encode sensory information and sub-serve behavior (Buzsáki and Chrobak, 1995; Buzsáki, 2008). This mechanism is particularly relevant to the olfactory system where spatiotemporal patterns of projection neuron (PN) activity act as robust markers of odor attributes (Laurent et al., 1996; Wehr and Laurent, 1996). In the insect antennal lobe (AL), a network of local inhibitory interneurons arborizes extensively throughout the AL (Leitch and Laurent, 1996) providing inhibitory input to the cholinergic PNs. Our theoretical work has attempted to elaborate the exact role of inhibition in the generation of odor specific PN responses (Bazhenov et al., 2001a,b; Assisi et al., 2011). In large-scale AL network models we characterized the inhibitory sub-network by its coloring (Assisi et al., 2011) and showed that it can entrain excitatory PNs to the odor specific patterns of transient synchronization. In this focused review, we further examine the dynamics of entrainment in more detail by simulating simple model networks in various parameter regimes. Our simulations in conjunction with earlier studies point to the key role played by lateral (between inhibitory interneurons) and feedback (from inhibitory interneurons to principal cells) inhibition in the generation of experimentally observed patterns of transient synchrony.

Keywords: inhibition, network, coloring, transient synchrony, olfaction, antennal lobe

INTRODUCTION

The dynamics of neuronal networks have often been associated with the detection of changes in the environment or as markers of expected and ongoing behavior. In olfactory networks this association is particularly well charted. Odor detection begins when odorant molecules bind to olfactory receptors, initiating a second messenger cascade that leads to the opening of ion channels, the depolarization (usually) of the receptor neuron cell membrane (Matthews and Reisert, 2003), and the generation of action potentials. Olfactory receptor neurons (ORNs) are preferentially sensitive to some odors and less so to others (Anderson et al., 1995; Hallem et al., 2004; Hallem and Carlson, 2006). This preference appears to be static, in that, odor rank orders are retained throughout the duration of the stimulus (Friedrich and Laurent, 2001; Bhandawat et al., 2007; Olsen et al., 2007). Chemically similar odors elicit responses from overlapping groups of ORNs that project onto a significantly smaller number of excitatory projection neurons (PNs) and inhibitory local interneurons (LNs) in the antennal lobe (AL). It has been suggested that the representation of a high dimensional, complex and as a first approximation, static input by convergence onto a few, albeit dynamic, PNs and LNs must enhance the fidelity of the input while unfolding the odor representation along the temporal dimension (Laurent, 2002) [In the locust ∼50,000 ORNs converge onto ∼900 PNs and 300 LNs (Ernst et al., 1977)]. In the AL, LNs extend extensive connections to PNs and to each other (Leitch and Laurent, 1993, 1996). The interaction between AL neurons produces spatiotemporal activity that evolves over multiple time scales (Laurent et al., 1996; Wehr and Laurent, 1996; Wilson et al., 2004). This evolution produces a progressive decrease in the overlap between the representations of chemically related odors (for example, several aromatic amino acids) (Friedrich and Laurent, 2001; Wilson et al., 2004) thereby increasing the ability of the system to discriminate between odors.

Local field potentials (LFP) recordings from different insects including locust (Laurent and Davidowitz, 1994; Laurent et al., 1996; Wehr and Laurent, 1996), moth (Ito et al., 2009), honeybee (Stopfer et al., 1997) and fly (Tanaka et al., 2009) have revealed fast 20–40 Hz oscillations that persist in spite of large changes in both the identity and the concentration of an odor (Stopfer et al., 2003). Successive oscillatory cycles are constructed from a dynamically evolving constellation of PN spikes. Odor attributes determine the identity of PNs that spike during a given oscillatory cycle. A given PN may be phase locked to specific cycles following the onset of an odor and spike randomly or remain silent during other cycles. Such transient synchronization between PNs was suggested to play a key role in encoding the odor representation (Laurent, 2002). Spatiotemporal patterns generated by PNs provide input to Kenyon cells (KCs) of the mushroom body via two pathways, a direct excitatory pathway from PNs to KCs and a feed-forward inhibitory pathway leading from lateral horn interneurons (LHIs) (Perez-Orive et al., 2004). Recent studies have also identified the role of a single neuron (the giant GABAergic cell) that provides feedback inhibitory input to KCs (Papadopoulou et al., 2011) (**Figure 1**).

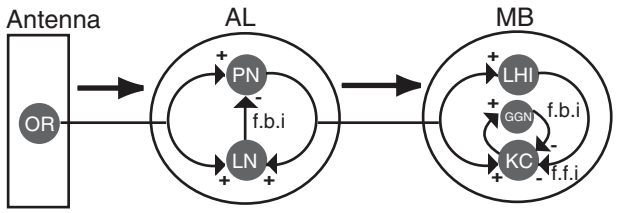

FIGURE 1 | Schematic of the insect olfactory system. Projection neurons (PNs) and local inhibitory interneurons (LNs) receive convergent input from odorant receptor (OR) neurons. Output from PNs is sent to Kenyon cells (KCs) of the Mushroom body (MB) and that also receive feed-forward inhibition from lateral horn interneurons (LHI). Giant GABAergic neuron (GGN) provides feedback inhibition to KCs.

In this article we discuss the relationship between the structure of the AL network and the dynamics that emerges from it, within the context of our earlier theoretical studies (Bazhenov et al., 2001a,b; Assisi et al., 2011). The efficacy and the relevance of the AL in representing an odor can only be inferred from its effect on subsequent layers, namely the KCs of the mushroom body. In contrast to PNs that fire promiscuously in response to an odor input, KCs generate a sparse and highly odor specific response. Owing to the intrinsic properties of KCs and network attributes (Perez-Orive et al., 2002, 2004), these neurons act as coincidence detectors of presynaptic input; KCs respond only to PN input that arrives during short windows of time bounded by the troughs of an oscillatory LFP cycle (Perez-Orive et al., 2002, 2004). If the same subset of PNs spike synchronously throughout an odor presentation, the odor would completely specified by identity of active PNs (an identity code) and the same postsynaptic KCs would be activated in every cycle. However, *in vivo* recordings from the locust AL show that different subsets of PNs are recruited and become transiently synchronized during different cycles of the LFP oscillation in a manner that is predictive of attributes of the odor. Thus, not only the identity of neurons that spike, but also the timing at which the spikes occur factor in encoding the odor representation (a spatiotemporal code).

Within the framework discussed here, the odor input is considered to be static and it stimulates the same set of PNs and LNs throughout the duration of the stimulus. This provides a useful first approximation to understand odor coding by the AL. Network interactions in the AL ensure that, different odors and odor attributes result in different spatiotemporal patterns of PN activity that, in turn, trigger different KC responses (Stopfer et al., 2003). Transient synchrony of AL neurons serves an important computational idea—it vastly expands the number of states available to encode odors when compared to a purely identity code. The upper bound on the number of states of the system increases as a power of the duration (measured in oscillatory cycles of the LFP) of the stimulus presentation. This could allow small differences in input to be amplified in a manner that it can be read out by KCs. While the animal need not, necessarily, take full advantage of this increased computational power, there is evidence to suggest that transient patterning plays a role in odor discrimination, particularly during difficult discrimination tasks (MacLeod et al., 1998).

Our goal is to elucidate the intrinsic and network mechanisms that are responsible for transient synchrony seen in the AL network. In particular we elaborate the specific role played by the inhibitory sub-network of the AL in facilitating the formation of transiently synchronous ensembles of PNs. In previous papers we have conjectured that the operational principles that govern the dynamics of the AL network include competition between LNs that generates epochs of time where specific LN sub-sets are activated. These, in turn, transiently entrain PNs that provide synchronous input to KCs (Bazhenov et al., 2001b; Assisi et al., 2011). This focused review further discusses and extends our theoretical ideas regarding odor coding in insect olfaction.

METHODS

Individual LNs and PNs were modeled by a single compartment with voltage and Ca^{2+} dependent currents described by Hodgkin-Huxley kinetics. The model LNs and PNs were constructed in order to closely emulate the experimentally observed dynamics of LNs and PNs using a parsimonious conductance based model [see (Bazhenov et al., 2001b) for a detailed description of the model parameters]. *In vivo*, PNs show overshooting sodium spikes in response to a constant depolarizing stimulus. LNs in contrast, generate low amplitude calcium spikelets that demonstrate spike frequency adaptation due to a Ca^{2+} dependent potassium current. Fast GABA (LN-PN and LN-LN connections) and nicotinic cholinergic synaptic currents (PN-LN connections) were modeled by first-order activation schemes. In addition to fast GABAergic synapses PNs in the AL are known to receive slow inhibitory input via $GABA_B$ type conductances (Bazhenov et al., 2001a; Wilson and Laurent, 2005). These slow responses tend to last over durations of 100 s of ms and generate long epochs when the neuron is hyperpolarized. However, since we were largely interested in how transient synchrony occurs on a fast time scale (lasting the duration of a few oscillatory cycles $<$100 ms) we chose to minimize the contributions of this slow form of inhibition. The specific connectivity of the networks and the excitatory and inhibitory coupling strengths were varied in individual simulations. The values are specified with individual figures. The model equations and other parameter values are specified in (Assisi et al., 2011). In the insect olfactory system, odor stimulation activates odor specific subset of PNs and LNs. In this study, to focus on the network driven spatiotemporal dynamics, we considered simplified stimulus model—a suprathreshold input was simultaneously provided to all PNs and LNs. The amplitude of this input was constant across all neurons except for a low amplitude additive noise term ($\sim$5% of the stimulus amplitude).

THE ROLE OF INHIBITION IN THE DYNAMICS OF SMALL NETWORKS

To understand the role of inhibition in generating transiently synchronous groups of PNs we first considered the dynamics of a reciprocally coupled pair of neurons consisting of an LN that inhibits a PN and receives excitatory input from it (**Figure 2A**). Both these neurons received identical depolarizing input that exceeded the spiking threshold of each neuron. The reciprocally coupled pair oscillated out of phase with a frequency of $\sim$25 Hz. The frequency of oscillation can vary as a function of the strength

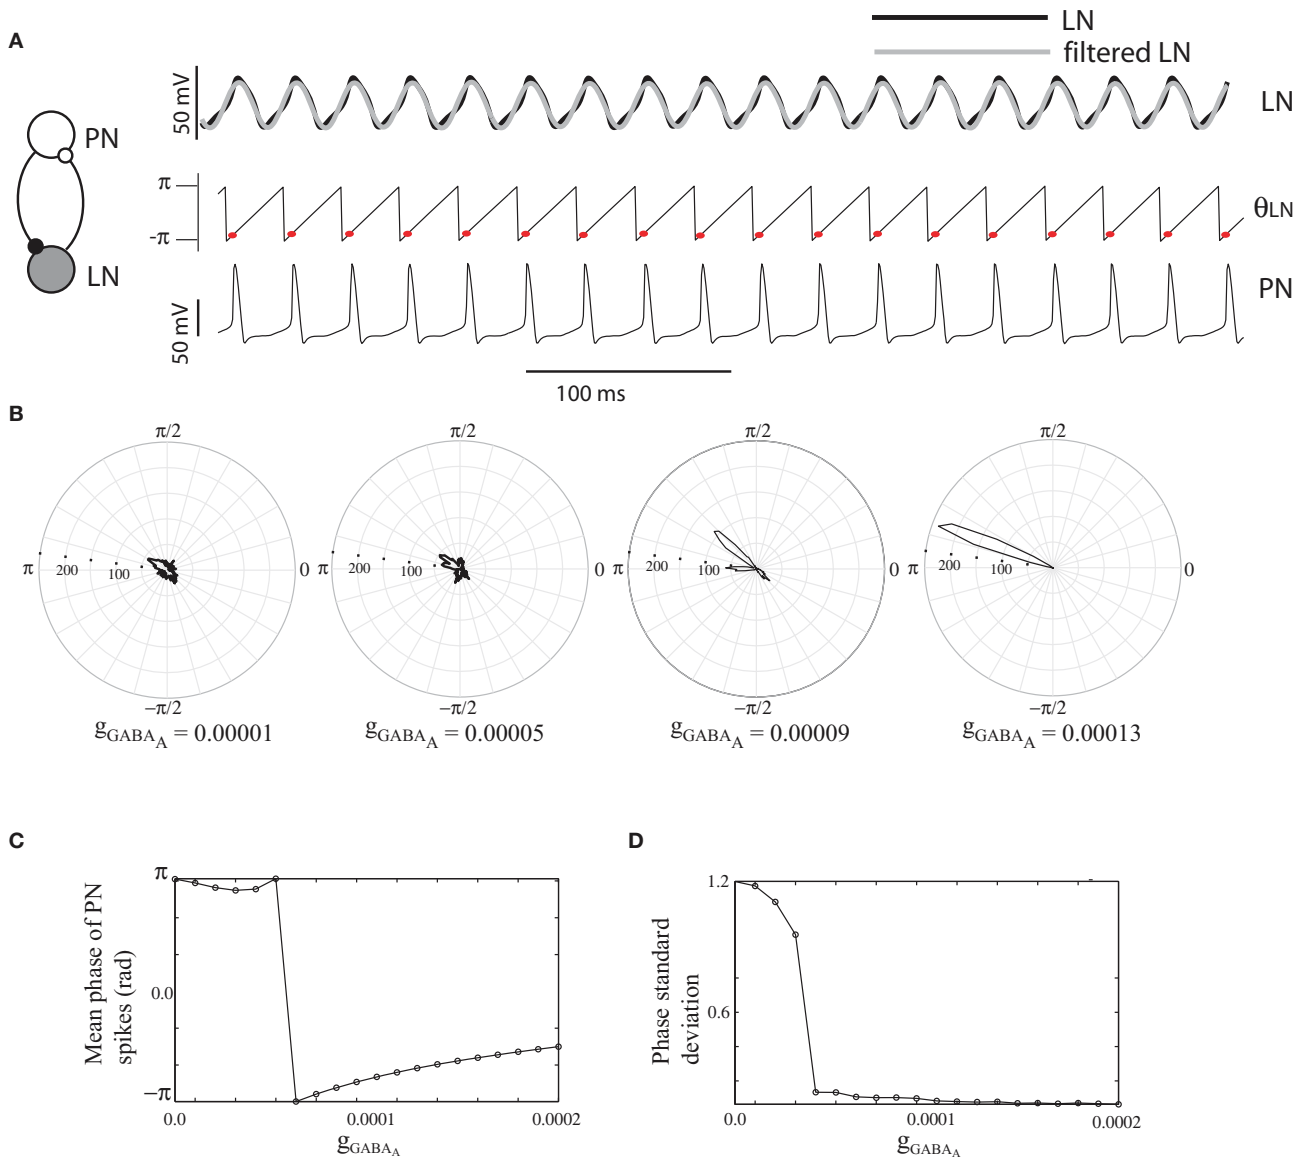

FIGURE 2 | (A) Single PN-LN pair. The top trace shows the membrane potential of the LN (black) and a frequency filtered version (gray trace, bandpass filter 10–40 Hz). The middle trace plots the instantaneous phase of the filtered LN membrane potential. The red dots show the value of the instantaneous phase at times when the PN generates a spike. The membrane potential of the PN is shown in the bottom trace. g_{ACh} (PN-LN) = 0.00001. **(B)** Phase distribution of PN spikes. The histogram shows the distribution of the phase at which the PN spikes for different values of the strength of inhibitory coupling from LN to PN (g_{GABA-A}). The distribution was generated over 20 s. **(C)** Mean phase of PN spikes. The mean phase at which the PN spikes is shown as a function of the inhibitory coupling strength. **(D)** The circular standard deviation of the phase as a function of the inhibitory coupling strength.

of the inhibitory coupling to the PN [see **Figure 2** in (Bazhenov et al., 2001b) for a detailed description]. We sought to determine the extent of PN entrainment by the inhibitory LN and the strength of inhibitory synaptic coupling required to entrain the PN. In order to quantify these features we measured the instantaneous phase of the LN by first filtering the LN membrane potential (**Figure 2A**, black line in the top trace) through a band pass filter (15–40 Hz) and calculating the Hilbert transform (**Figure 2A**, middle trace) of the resulting smoothed trace (**Figure 2A**, gray line in the top trace). The instantaneous phase thus determined was sampled at times when the PN generated a spike (red circles in **Figure 2A**, middle trace). We simulated this simple network over 20 s and calculated the distribution of these phase points corresponding to the times when the PN generated a spike (**Figure 2B**). The spread of the distribution provided a measure of the degree to which the LN entrained the dynamics of the PN (**Figure 2D**). As we varied the strength of the inhibitory synaptic coupling, the circular standard deviation of the phase of PN spikes with respect to the LN oscillation decreased until, at a specific value of g_{GABA-A}, it abruptly dropped from

∼1.0–∼0.05 radians followed by a gradual monotonic decrease (**Figure 2D**). The effect of an LN spike was to shift the onset of the following PN spike to a later phase of the oscillation as the coupling strength increased (**Figure 2C**). The mean phase of the distribution moved toward the peak of the LN oscillatory cycle. Thus individual PNs can be entrained by its coupling to a single LN if the strength of the inhibitory input exceeds a specific threshold.

Next we sought to determine the collective dynamics of multiple PNs, all receiving identical inhibitory input from a single LN. We simulated a network of 10 PNs that were inhibited by a single LN and provided excitatory input to it (**Figure 3A**).

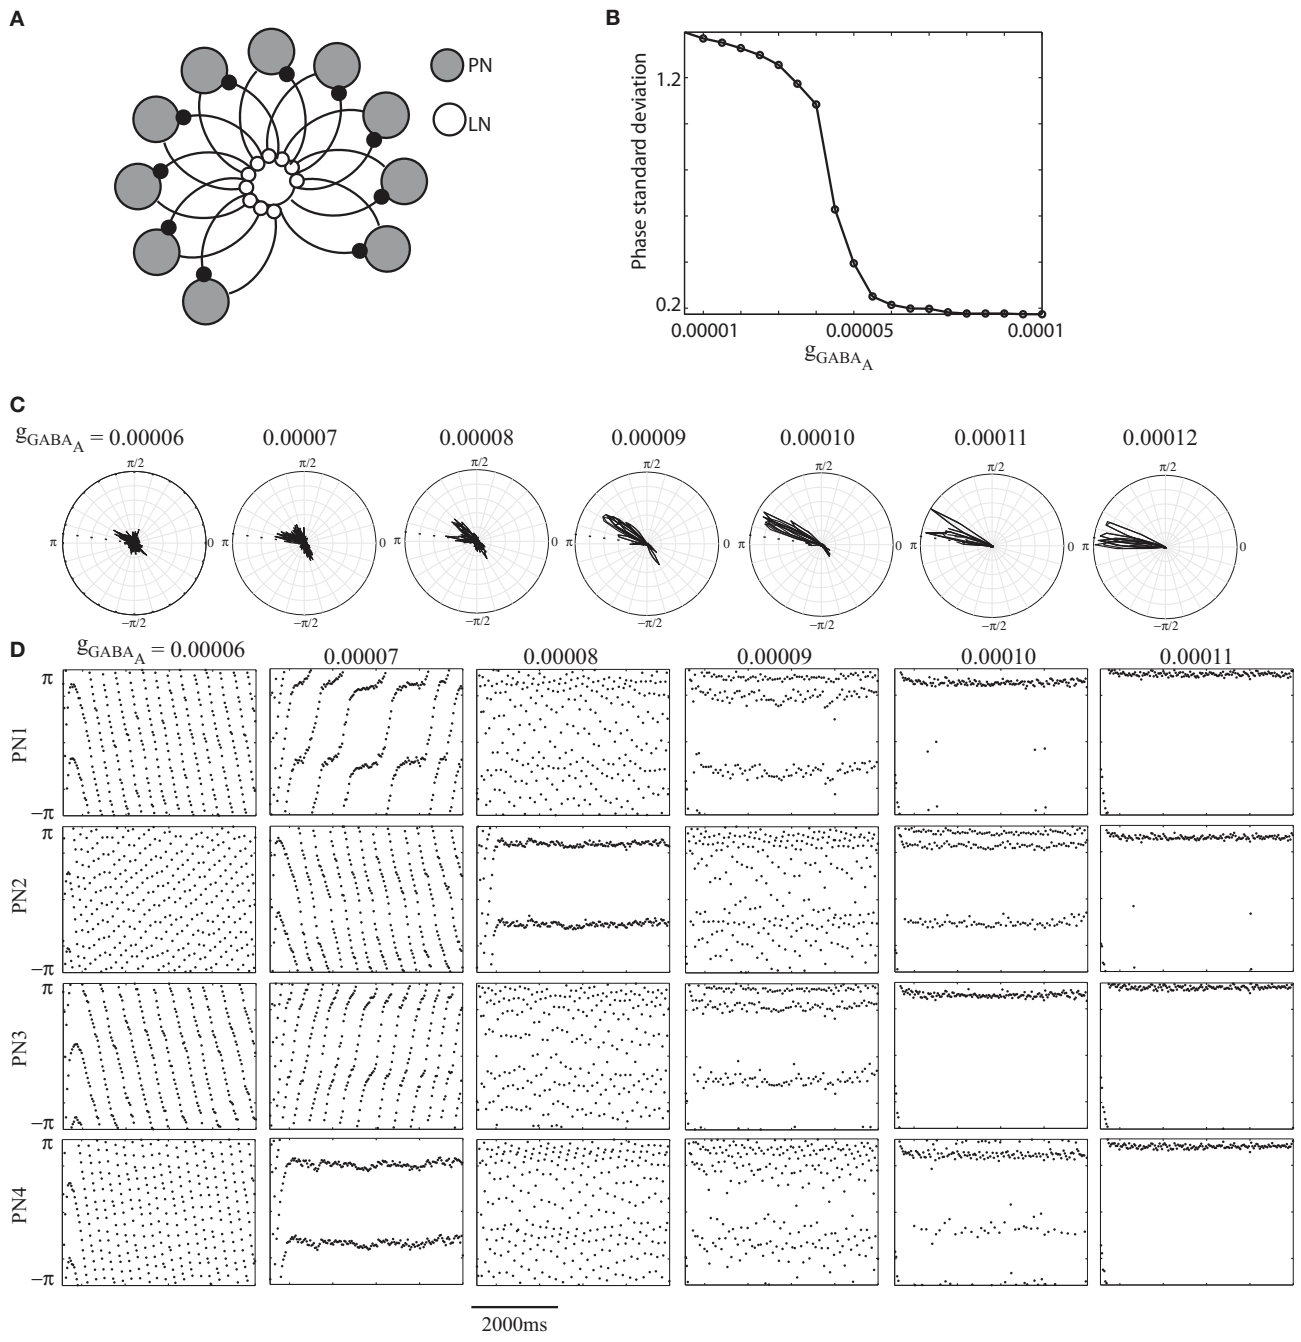

FIGURE 3 | (A) Network consisting of a single LN inhibiting 10 PNs. The LN receives excitatory input from all PNs. g_{ACh} (PN-LN) = 0.00001, g_{GABA-A} (LN-LN) = 0.00001. **(B)** The standard deviation of the PN spike phases as a function increasing strength of inhibitory coupling. **(C)** The distribution of the phase of PN spikes for different values of g_{GABA-A}. **(D)** PN dynamics. Each dot in a panel corresponds to the instantaneous phase of a PN spike at a given time. The phase dynamics of 4 PNs of the 10 PNs simulated is shown. Each column of panels corresponds to a different value of $g_{GABA-A(LN-PN)}$.

Here too we found that the collective dynamics of all PNs was strongly entrained by the LN oscillations when the strength of the inhibitory coupling exceeded a threshold value (**Figures 3B,C**). The standard deviation of the phase at which the PNs generated spikes abruptly decreased as the strength of inhibitory coupling increased. The dynamics of the PNs however showed complex patterns of activity relative to the LN (**Figure 3D**). For different values of the coupling strength, a given PN could show phase drift relative to the LN oscillations, multi-frequency phase locking and, for the strongest coupling strengths, 1:1 phase locked dynamics (**Figure 3D**).

These simulations show that LNs can effectively entrain the dynamics of PNs. Therefore, we hypothesize that transient synchrony between PNs, as seen in experimental recordings in the locust AL, must, to a significant extent, be driven by its interaction with the LN sub-network. Evidence for this assertion also comes from experiments where the influence of fast inhibition mediated by $GABA_A$ was selectively abolished by the application of $GABA_A$ antagonist picrotoxin (Stopfer et al., 1997). This had the effect of curtailing a prominent 20 Hz oscillation that is ubiquitous during odor stimulation. Since the experimentally measured oscillatory local field potential can be viewed as the synchronized activity of ensembles of PNs, we can, as subsequent modeling studies have (Bazhenov et al., 2001b), infer that inhibition indeed entrains PN activity. Our simulations (**Figures 2** and **3**) examine the extent of this entrainment over a broader parameter range than that simulated in earlier studies (Bazhenov et al., 2001b; Assisi et al., 2011). Furthermore, reciprocal inhibition between LNs can enforce competition between groups of inhibitory neurons. This would lead to some groups of LNs being activated in response to an odor, while other groups are quiescent or hyperpolarized by the active LNs (Assisi et al., 2011). Transient synchrony within groups of LNs, coupled with their ability to drive the activity of PNs provides a potential mechanism that explains the observed spatiotemporal patterning of PNs in the AL (see more on this below).

The role of inhibition in synchronizing the activity of excitatory PNs has also been examined models of Limax (Ermentrout et al., 1998) and honeybee (Linster and Cleland, 2001). Synchronization driven by inhibition is not limited to olfactory networks. In the hippocampus, GABAergic neurons form the hubs of a network that synchronizes the activity of pyramidal cells (Bonifazi et al., 2009). Feedback inhibition mediated by LNs can effectively synchronize distributed pyramidal cells, a mechanism termed pyramidal interneuronal network gamma (PING) (Borgers and Kopell, 2003, 2005; Borgers et al., 2005). This mechanism has also been implicated in the generation of gamma band synchrony (Llinas and Ribary, 1993; Singer and Gray, 1995; Wang and Buzsaki, 1996).

RECIPROCAL INHIBITION CAN BE RELATED TO THE COLORING OF THE NETWORK

The two LNs in **Figure 4** reciprocally inhibit each other. Over a wide range of parameters these neurons would tend to spike out of phase [see (Vreeswijk et al., 1994) for exceptions]. When one of the neurons generated a burst of spikes it suppressed activity in the other neuron. In the network simulated above, which neuron spikes is determined by the calcium concentration in the cell. Increased intracellular Ca^{2+} signals a lower propensity for firing (see (Ahn et al., 2010) for a detailed view of the effects of Ca^{2+} on the spiking pattern). We have shown that burst alternation was determined by the activation of Ca^{2+}-dependent K^+ currents (Assisi et al., 2011). Each spike led to an increase in the Ca^{2+} concentration within a cell and additional K (Ca^{2+}) activation that consequently delayed the onset of the following spike, a phenomenon known as spike frequency adaptation. When the frequency of spiking reduced below certain threshold, the inhibition provided by this neuron was not sufficient to keep the other neuron from spiking; as a result, the quiescent neuron switched to an active state and suppressed activity in the post-synaptic neuron. Antagonistic interactions between reciprocally coupled inhibitory neurons allowed us to derive a relationship between the dynamics of neurons and a structural characteristic of the underlying network, namely, its coloring.

A coloring of the network is a prescription that assigns different colors to nodes (neurons) that are directly connected to each other (Chartrand, 1984). The minimum number of colors required to color a network is known as its chromatic number. In the example above, the two neurons are reciprocally coupled and therefore must be assigned different colors. If the coupling between neurons is inhibitory, we anticipate that neurons associated with different colors, upon stimulation by the

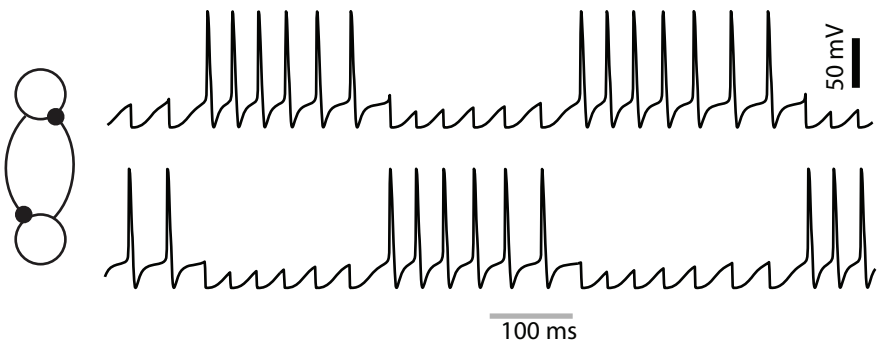

FIGURE 4 | A pair of reciprocally coupled LNs generate alternating patterns of spikes. The two traces correspond to different LNs shown in the figure.

same depolarizing input, will compete with each other and generate spikes (or bursts of spikes) at different times. In contrast, neurons that do not compete (are associated with the same color) will spike in approximate synchrony. This relationship between the coloring and the dynamics of the inhibitory network is clearly evident in the simple network constructed in **Figure 4**. However, does this relationship generalize to larger and more complex networks? To address this question we constructed networks that possessed the required coloring [see (Assisi et al., 2011)]. For example to construct a network with two colors, we generated two groups of neurons that were all-to-all connected across groups but possessed no connections within groups. **Figure 5A1** shows an example of one such network consisting of 40 neurons with 20 being associated with one color and the remaining with another

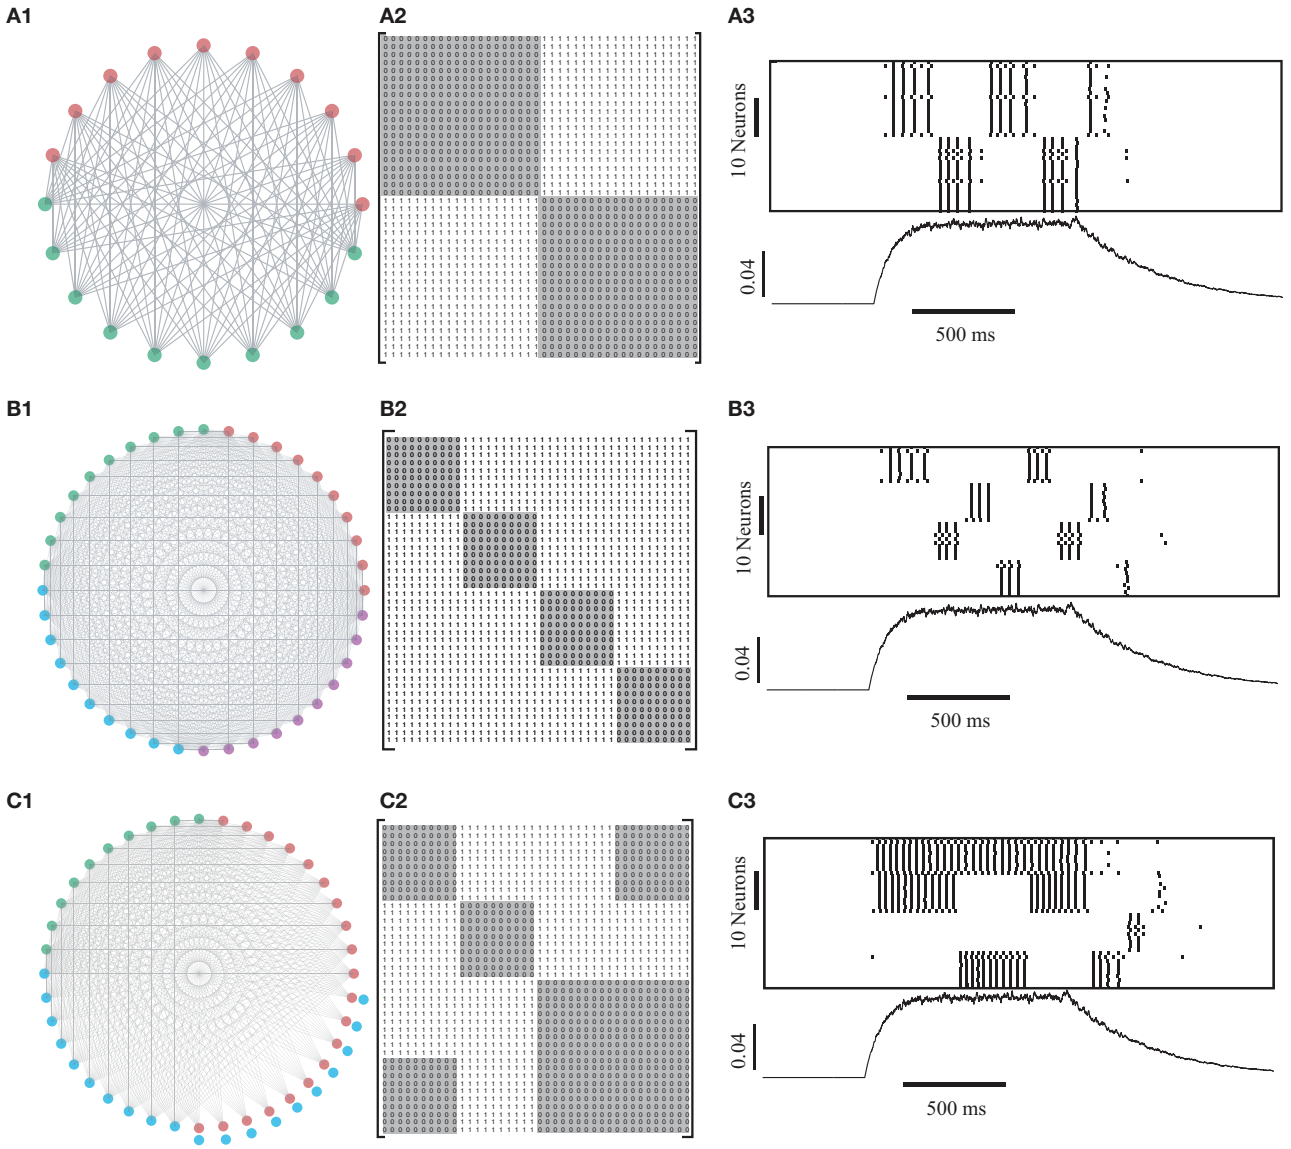

FIGURE 5 | Dynamics of a network of reciprocally connected inhibitory LNs. (A) Bipartite graph with chromatic number 2. Each LN is reciprocally connected to all the neurons that are associated with a color different from its own **(A1)**. The adjacency matrix of the network consists of diagonal blocks of zeros with ones elsewhere **(A2)**. Raster plot showing the activity of 40 neurons in response to identical input to all the neurons. LNs associated with the same color spike in approximate synchrony **(A3)**. The traces below the raster plots in Aiii, Biii, Ciii show the input that is given to all the inhibitory neurons. The parameter values used for this simulation were g_{GABA-A} (LN-LN) = 0.00007. **(B)** A graph with four colors can be constructed by picking four groups of neurons that do not extend connections within a group but possess all to all reciprocal inhibitory connections with other groups **(B1)**. The adjacency matrix consists of four diagonal blocks of zeros with ones elsewhere **(B2)**. As in the previous case neurons associated with the same color tend to spike in synchrony **(B3)** (g_{GABA-A} (LN-LN) = 0.00007). **(C)** Graph with multiple colorings. Three groups of neurons possess all-to-all reciprocal connections between groups. A fourth group of neurons extend connections only to the group marked in green and not to the others. Therefore, it may be associated with the blue or the red group **(C1)**. Neurons associated with this group tend to spike in synchrony with either the red or the blue group but remain silent when the green groups spikes **(C3)** (g_{GABA-A} (LN-LN) = 0.00002). The corresponding adjacency matrix is shown in **C2**.

color. The adjacency matrix of this network consisted of two 20×20 diagonal blocks of zeros and ones in all other locations (**Figure 5A2**). The diagonal blocks of zeros are the shaded regions of the adjacency matrices shown in **Figure 5A2**. The dynamics of this network with chromatic number, two, clearly showed that neurons associated with different colors generated an alternating pattern of activity (**Figure 5A3**). The antagonistic interactions between inhibitory groups persisted even as the number of colors increased (**Figure 5B**).

A number of experimental recordings show that PNs, and therefore LNs that entrain them, form transiently synchronous groups over the duration of an odor. A given PN may participate in more than one synchronous group. In **Figures 5A** and **5B**, however, we show that a given LN associated with a particular color, can only spike in synchrony with LNs that share the same color and not with any other group. If LNs indeed drive the synchrony of PN spikes (see below), this would imply that the constitution of every synchronously spiking group of PNs would never change. One can circumvent this issue by constructing networks that possess multiple colorings. For example, the network in **Figure 5C1** has chromatic number three. However, in contrast with the networks shown in **Figures 5A** and **5B**, this network (**Figure 5C1**) can be colored using three colors in two different ways. A group of ten neurons in **Figure 5C1** can be colored either red or blue but not green since these neurons are connected to the green group but not to the other two. The adjacency matrix of this network has three 10×10 diagonal blocks of zeros in addition to some off-diagonal blocks of zeros (**Figure 5C2**) [see **Figure 3** in (Assisi et al., 2011) for a simplified rendering of this graph where all neurons associated with a color are grouped into a single node]. The dynamics of this network under constant and identical deploarizing input to all neurons evolves such that the group of neurons that can be colored using multiple colors spikes in synchrony with both the red and the blue groups while falling silent when the green group of neurons spike. Multiple colorings can therefore be associated with transient synchrony where individual neurons can participate in multiple synchronous groups as seen in experiments. This example also shows that the neurons in the inhibitory network do not have to be directly connected to each other to spike asynchronously.

While our simulations show that antagonistic interactions between LNs generates clusters of synchronously spiking neurons that respect the coloring of the network, this need not always be the case. In all-to-all connected (Golomb and Rinzel, 1993; Wang and Buzsaki, 1996) and random networks (Wang and Buzsaki, 1996), the dynamics can depend on the interaction of intrinsic (spike afterhyperpolarization, heterogeneities in spiking frequency) and synaptic (reversal potential, synaptic current decay time, sparseness of connections) factors. Clustered solutions of the type observed here form a sub-set of the full dynamical repertoire of purely inhibitory networks. Heterogeneities in network structure [**Figure 6** in (Assisi et al., 2011)] and dense connectivity between neurons, resulting in a larger number of colors [**Figure 2** in (Assisi et al., 2011)] can perturb the coloring based dynamics of the system. However, the inclusion of feedback excitation has been shown to reduce the extent of this perturbation [**Figure 4** in(Assisi et al., 2011)].

TRANSIENT SYNCHRONY IN PROJECTION NEURONS IS DRIVEN BY THE INHIBITORY SUB-NETWORK

In our simulations we showed that feedback GABAergic inhibition mediated by LNs could entrain PNs to spike in synchrony (**Figures 2** and **3**). Further, reciprocally coupled networks of LNs could, as determined by the coloring of the network, form transiently synchronous groups that moved in and out of synchrony with each other over the duration of the odor input (**Figures 4** and **5**). In what follows we show that these groups of LNs can entrain PNs to generate the kind of transiently synchronous dynamics seen in recordings from PNs in the locust AL. This entrainment is seen in PNs regardless of whether the LN networks are uniquely colored or possess multiple colorings. We constructed a network consisting of forty PNs and forty LNs. The LN sub-network consisted of four groups (colors) of neurons that were all-to-all connected across groups but received no connections from within each group (**Figure 5B**). Each LN group provided inhibitory input to one of four groups of PNs but received excitatory input from all the PNs (**Figure 6A**). Non-oscillatory external input (**Figure 6B**, bottom panel) triggered $\sim$30 Hz non-synchronous spiking in all PNs of the network (**Figure 6B**, middle panel). As expected, the LNs generated an alternating pattern of activity consistent with the coloring of the underlying network (**Figure 6B**, top panel). Each burst of spikes by an active LN group synchronized the activity of post-synaptic PNs. As a result, all PNs that received input from an active LN group showed synchronized $\sim$20 Hz oscillations (**Figure 6B**, middle panel). When the group of LNs became quiescent the post-synaptic PNs drifted from synchrony during subsequent epochs of time when other LN groups generated spikes. These dynamics lead to a complex pattern of synchronization where different groups of PNs became synchronized during different epochs of odor stimulation.

DISCUSSION

We demonstrated that competitive interactions between LNs can lead to the formation of groups of synchronously spiking LNs that can effectively entrain the dynamics of PNs. The identity of LNs that participate in each group is determined by a characteristic feature of the connectivity structure of the inhibitory network, namely, its coloring. Our simulations in **Figure 6** demonstrate that spikes generated by PNs are phase locked to the dynamics of those LNs that provide direct inhibitory input to it. The degree of PN synchronization is dependent on the strength of the inhibitory coupling from LNs. In simple networks consisting of a single inhibitory neuron that was reciprocally connected to multiple PNs, we demonstrated multi-frequency coordination between PNs and the LN. Multifrequency coordination has also been reported in Lotka-Volterra models with competing interactions between the nodes (Rabinovich et al., 2006). Such dynamics clearly defines a widening of the dynamical repertoire of that these networks are capable of.

The critical feature of out model is that the strongest inhibition determines the identities of PNs that "talk" to the downstream neurons in the mushroom bodies. Indeed only PNs that receive sufficiently strong inhibition can (transiently) synchronize and drive responses in the KCs that operate as spike coincidence detectors. Therefore, we can classify our model as multiplicative—the

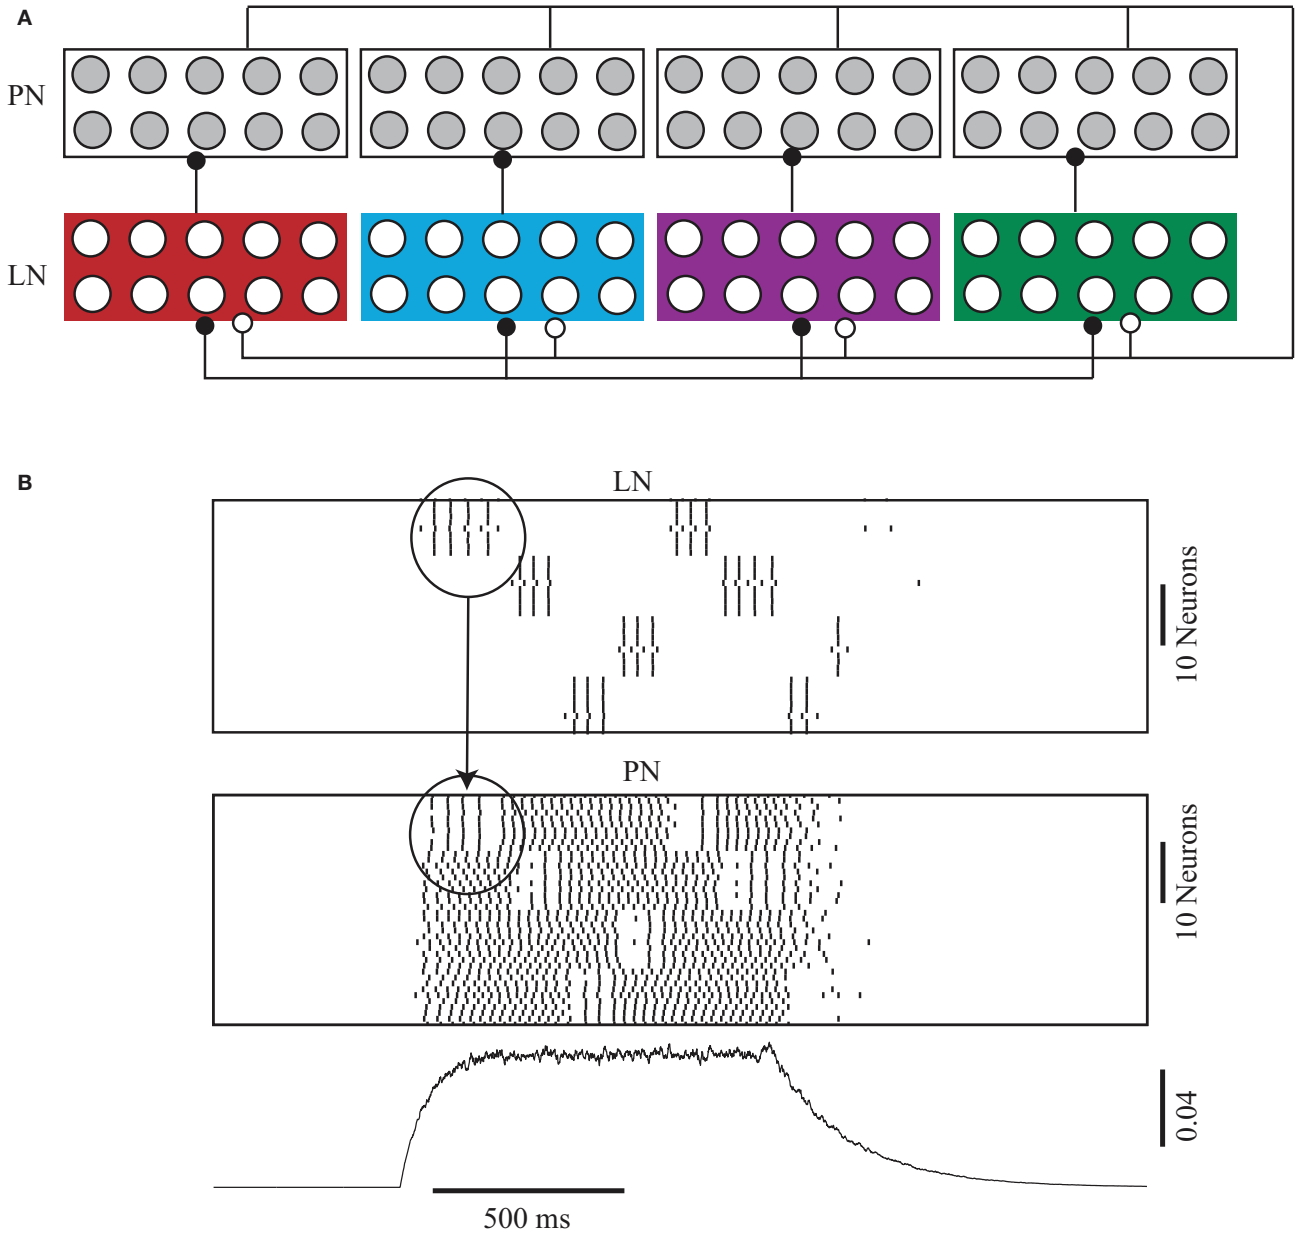

FIGURE 6 | (A) A network of LNs with chromatic number 4 is connected to a population of PNs. The PNs were separated into four groups. Each group received inhibitory input only from LNs that were associated with one color and not others. The PNs provided excitatory input to all LNs irrespective of the coloring. (g_{GABA-A} (LN-LN) = 0.00004, g_{GABA-A} (LN-PN) = 0.00006, g_{Ach} (PN-LN) = 0.00004). **(B)** PNs (top raster) that received inhibitory input from LNs (bottom raster) of a particular color tended to spike synchronously only when those LNs were activated. The black trace at the bottom shows the input provided to all PNs and LNs.

input from olfactory receptor neurons to PNs is multiplied by feedback inhibition within AL to determine which PNs transmit information to the mushroom body. PNs receiving sensory input but no inhibitory feedback cannot reach the mushroom bodies. This is in contrast to some other theoretical models of olfactory processing (in vertebrates) where inhibition balances excitation in many principal neurons of the olfactory bulb and only the "error" signal is transmitted to the subsequent level—piriform cortex (Koulakov and Rinberg, 2011).

A long-standing conjecture in neuroscience has been that neuronal networks possess mechanisms that can string groups of synchronously spiking neurons into elaborate temporal sequences (Hebb, 1949). These spatiotemporal patterns, termed phase sequences, then form the basis of all perception (Gray and Singer, 1989), memory (Cheng and Frank, 2008) and action (Bouyer et al., 1987). Here we demonstrated that the coloring of the network determines the identity of neurons that spike together. However, the coloring does not place any specific constrains

on the sequence in which successive groups of neurons are recruited. The occurrence of a specific sequence of spiking is influenced by the nature of the input, the history, and the current state of the network. In the inhibitory networks simulated here the intracellular Ca^{2+} concentration within a cell plays a prominent role in determining the order in which groups of neurons generate spikes. The lowest Ca^{2+} concentration determined the minimal activation of Ca^{2+} dependent K^{+} currents and, therefore, the most excitable group of LNs (Ahn et al., 2010; Assisi et al., 2011). In all the simulations in this paper we introduced a small amount of variability in the intrinsic properties across PNs and LNs to ensure that the network dynamics does not settle into a pathological and unstable state. This variability could potentially lead to asymmetries that define a specific ordering of the PN sequences. Explicitly imposed asymmetries in the properties of the neurons such as the excitability of different groups (either intrinsic or stimulus-dependent) could be used to enforce the stability of some sequences in lieu of myriad others that could also be present in the completely symmetric network. Finally, the LN networks simulated here possessed reciprocal connections between pairs. Directional connectivity between neurons is known to introduce a strong asymmetry in the network and restricts number of stable sequences of activity. This can be used to construct networks that possess a priori specified patterns of activity.

Our study revealed a link between the structure of the excitatory-inhibitory network and its dynamics. It presents a new approach to understand how the dynamics of the network can be predicted from the connectivity and may lead to strategies that could allow us to infer the connectivity of the network from its observed dynamics.

ACKNOWLEDGMENT

This research is supported by grant from NIH-NIDCD.

REFERENCES

Ahn, S., Smith, B. H., Borisyuk, A., and Terman, D. (2010). Analyzing neuronal networks using discrete-time dynamics. *Physica D* 239, 515–528.

Anderson, P., Hansson, B. S., and Lofquist, J. (1995). Plant-odour-specific receptor neurones on the antennae of female and male spodoptera littoralis. *Physiol. Entomol.* 20, 189–198.

Assisi, C., Stopfer, M., and Bazhenov, M. (2011). Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning. *Neuron* 69, 373–386.

Bazhenov, M., Stopfer, M., Rabinovich, M., Abarbanel, H. D., Sejnowski, T. J., and Laurent, G. (2001a). Model of cellular and network mechanisms for odor-evoked temporal patterning in the locust antennal lobe. *Neuron* 30, 569–581.

Bazhenov, M., Stopfer, M., Rabinovich, M., Huerta, R., Abarbanel, H. D., Sejnowski, T. J., and Laurent, G. (2001b). Model of transient oscillatory synchronization in the locust antennal lobe. *Neuron* 30, 553–567.

Bhandawat, V., Olsen, S. R., Gouwens, N. W., Schlief, M. L., and Wilson, R. I. (2007). Sensory processing in the Drosophila antennal lobe increases reliability and separability of ensemble odor representations. *Nat. Neurosci.* 10, 1474–1482.

Bonifazi, P., Goldin, M., Picardo, M. A., Jorquera, I., Cattani, A., Bianconi, G., Represa, A., Ben-Ari, Y., and Cossart, R. (2009). GABAergic hub neurons orchestrate synchrony in developing hippocampal networks. *Science.* 326, 1419–1424.

Borgers, C., Epstein, S., and Kopell, N. J. (2005). Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc. Natl. Acad. Sci. U.S.A.* 102, 7002–7007.

Borgers, C., and Kopell, N. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Comput.* 15, 509–538.

Borgers, C., and Kopell, N. (2005). Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons. *Neural Comput.* 17, 557–608.

Bouyer, J. J., Montaron, M. F., Vahnee, J. M., Albert, M. P., and Rougeul, A. (1987). Anatomical localization of cortical beta rhythms in cat. *Neuroscience* 22, 863–869.

Buzsáki, G. (2008). Neural inhibition. *Scholarpedia* 9.

Buzsáki, G., and Chrobak, J. J. (1995). Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr. Opin. Neurobiol.* 5, 504–510.

Chartrand, G. (1984). *Introductory Graph Theory.* New York, NY: Dover Publications.

Cheng, S., and Frank, L. M. (2008). New experiences enhance coordinated neural activity in the hippocampus. *Neuron* 57, 303–313.

Ermentrout, B., Flores, J., and Gelperin, A. (1998). Minimal model of oscillations and waves in the limax olfactory lobe with tests of the model's predictive power. *J. Neurophysiol.* 79, 2677–2689.

Ernst, K. D., Boeckh, J., and Boeckh, V. (1977). A neuroanatomical study on the organization of the central antennal pathways in insects. *Cell Tissue Res.* 176, 285–306.

Friedrich, R. W., and Laurent, G. (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science* 291, 889–894.

Golomb, D., and Rinzel, J. (1993). Dynamics of globally coupled inhibitory neurons with heterogeneity. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* 48, 4810–4814.

Gray, C. M., and Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 86, 1698–1702.

Hallem, E. A., and Carlson, J. R. (2006). Coding of odors by a receptor repertoire. *Cell* 125, 143–160.

Hallem, E. A., Ho, M. G., and Carlson, J. R. (2004). The molecular basis of odor coding in the Drosophila antenna. *Cell* 117, 965–979.

Hebb, D. O. (1949). *The Organization of Behavior.* New York, NY: John Wiley and Sons.

Ito, I., Bazhenov, M., Ong, R. C., Raman, B., and Stopfer, M. (2009). Frequency transitions in odor-evoked neural oscillations. *Neuron* 64, 692–706.

Koulakov, A. A., and Rinberg, D. (2011). Sparse incomplete representations: a potential role of olfactory granule cells. *Neuron* 72, 124–136.

Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nat. Rev. Neurosci.* 3, 884–895.

Laurent, G., and Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science* 265, 1872.

Laurent, G., Wehr, M., and Davidowitz, H. (1996). Temporal representations of odors in an olfactory network. *J. Neurosci.* 16, 3837–3847.

Leitch, B., and Laurent, G. (1993). Distribution of GABAergic synaptic terminals on the dendrites of locust spiking local interneurones. *J. Comp. Neurol.* 337, 461–470.

Leitch, B., and Laurent, G. (1996). GABAergic synapses in the antennal lobe and mushroom body of the locust olfactory system. *J. Comp. Neurol.* 372, 487–514.

Linster, C., and Cleland, T. A. (2001). How spike synchronization among olfactory neurons can contribute to sensory discrimination. *J. Comput. Neurosci.* 10, 187–193.

Llinas, R., and Ribary, U. (1993). Coherent 40-Hz oscillation characterizes dream state in humans. *Proc. Natl. Acad. Sci. U.S.A.* 90, 2078–2081.

MacLeod, K., Backer, A., and Laurent, G. (1998). Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 395, 693–698.

Matthews, H. R., and Reisert, J. (2003). Calcium, the two-faced messenger of olfactory transduction and adaptation. *Curr. Opin. Neurobiol.* 13, 469–475.

Olsen, S. R., Bhandawat, V., and Wilson, R. I. (2007). Excitatory interactions between olfactory processing channels in the Drosophila antennal lobe. *Neuron* 54, 89–103.

Papadopoulou, M., Cassenaer, S., Nowotny, T., and Laurent, G. (2011). Normalization for sparse encoding of odors by a wide-field interneuron. *Science* 332, 721–725.

Perez-Orive, J., Bazhenov, M., and Laurent, G. (2004). Intrinsic and

circuit properties favor coincidence detection for decoding oscillatory input. *J. Neurosci.* 24, 6037–6047.

Perez-Orive, J., Mazor, O., Turner, G. C., and Cassenaer, S. (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science*. 297, 359–365.

Rabinovich, M., Huerta, R., and Varona, P. (2006). Heteroclinic synchronization: ultrasubharmonic locking. *Phys. Rev. Lett.* 96, 014101.

Singer, W., and Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* 18, 555–586.

Stopfer, M., Bhagavan, S., Smith, B. H., and Laurent, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature*. 390, 70–74.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.

Tanaka, N. K., Ito, K., and Stopfer, M. (2009). Odor-evoked neural oscillations in Drosophila are mediated by widely branching interneurons. *J. Neurosci.* 29, 8595–8603.

Vreeswijk, C., Abbott, L., and Bard Ermentrout, G. (1994). When inhibition not excitation synchronizes neural firing. *J. Comput. Neurosci.* 1, 313–321.

Wang, X. J., and Buzsaki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J. Neurosci.* 16, 6402–6413.

Wehr, M., and Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384, 162–166.

Wilson, R. I., and Laurent, G. (2005). Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the Drosophila antennal lobe. *J. Neurosci.* 25, 9069–9079.

Wilson, R. I., Turner, G. C., and Laurent, G. (2004). Transformation of olfactory representations in the Drosophila antennal lobe. *Science* 303, 366–370.

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.

Received: 14 October 2011; paper pending published: 09 December 2011; accepted: 03 April 2012; published online: 18 April 2012.

*Citation: Assisi C and Bazhenov M (2012) Synaptic inhibition controls transient oscillatory synchronization in a model of the insect olfactory system. Front. Neuroeng. **5**:7. doi: 10.3389/fneng.2012.00007*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 25 October 2011
doi: 10.3389/fneng.2011.00012

Detailed characterization of local field potential oscillations and their relationship to spike timing in the antennal lobe of the moth *Manduca sexta*

***Kevin C. Daly*[1]*, *Roberto F. Galán*[2], *Oakland J. Peters*[1] and *Erich M. Staudacher*[1]**

[1] *Department of Biology, West Virginia University, Morgantown, WV, USA*
[2] *Department of Neurosciences, Case Western Reserve, Cleveland, OH, USA*

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Vassiliy Tsytsarev, George Mason University, USA
Paul Szyszka, Universität Konstanz, Germany

***Correspondence:**
Kevin C. Daly, Department of Biology, West Virginia University, PO Box 6057, Morgantown, WV 26506, USA.
e-mail: kevin.daly@mail.wvu.edu

The transient oscillatory model of odor identity encoding seeks to explain how odorants with spatially overlapped patterns of input into primary olfactory networks can be discriminated. This model provides several testable predictions about the distributed nature of network oscillations and how they control spike timing. To test these predictions, 16 channel electrode arrays were placed within the antennal lobe (AL) of the moth *Manduca sexta*. Unitary spiking and multi site local field potential (LFP) recordings were made during spontaneous activity and in response to repeated presentations of an odor panel. We quantified oscillatory frequency, cross correlations between LFP recording sites, and spike–LFP phase relationships. We show that odor-driven AL oscillations in *Manduca* are frequency modulating (FM) from ~100 to 30 Hz; this was odorant and stimulus duration dependent. FM oscillatory responses were localized to one or two recording sites suggesting a localized (perhaps glomerular) not distributed source. LFP cross correlations further demonstrated that only a small ($r < 0.05$) distributed and oscillatory component was present. Cross spectral density analysis demonstrated the frequency of these weakly distributed oscillations was state dependent (spontaneous activity = 25–55 Hz; odor-driven = 55–85 Hz). Surprisingly, vector strength analysis indicated that unitary phase locking of spikes to the LFP was strongest during spontaneous activity and dropped significantly during responses. Application of bicuculline, a $GABA_A$ receptor antagonist, significantly lowered the frequency content of odor-driven distributed oscillatory activity. Bicuculline significantly reduced spike phase locking generally, but the ubiquitous pattern of increased phase locking during spontaneous activity persisted. Collectively, these results indicate that oscillations perform poorly as a stimulus-mediated spike synchronizing mechanism for *Manduca* and hence are incongruent with the transient oscillatory model.

Keywords: olfaction, odor coding, oscillations, synchrony, $GABA_A$, olfactory bulb, antennal lobe

INTRODUCTION

Odor stimulation drives a combinatorial spatio-temporal response in first-order olfactory processing centers, the insect antennal lobe (AL), and the vertebrate olfactory bulb (OB). Several aspects of these physiological responses have been correlated with an animal's ability to discriminate between even subtly different odorant stimuli in behavioral assays (Stopfer et al., 1997; Daly et al., 2001; Linster et al., 2001; Wright et al., 2002; Galán et al., 2004). For example, odor can drive local field potential oscillations (LFPOs) at species-specific frequencies (Gelperin and Tank, 1990; Heinbockel et al., 1998). Superimposed on these LFPOs are what has been described as slower temporal patterns of action potentials among populations of principal output cells of the AL/OB (for review see Laurent et al., 2001). At least a subset of action potentials from these responses has been shown to correlate to a specific phase range (i.e., phase lock) of the LFPOs in the locust (Laurent and Davidowitz, 1994; Laurent and Naraghi, 1994; Laurent et al., 1996a,b, 2001; MacLeod and Laurent, 1996; Wehr and Laurent, 1996).

Selective disruption of $GABA_A$ receptor signaling in the locust AL interferes with fast inhibitory synaptic interactions, thereby disrupting odor-driven oscillations in the AL and its primary projection fields (MacLeod and Laurent, 1996). Disruption of $GABA_A$ mediated oscillations also appears to affect discrimination of closely related monomolecular odors in mouse and honeybee (Stopfer et al., 1997; Nusser et al., 2001). Thus, based on comparative data from locust (Laurent and Davidowitz, 1994; Laurent and Naraghi, 1994; Wehr and Laurent, 1996, 1999), honeybee (Stopfer et al., 1997), zebra fish (Friedrich and Laurent, 2001), and mice (Nusser et al., 2001), a general model of olfactory encoding has been proposed, positing that stimulus driven LFPOs synchronize transiently active spiking responses across a distributed network of principal output cells in order to mediate enhanced odor discrimination over time in downstream brain centers, such as the insect mushroom bodies (MB; for review see Laurent, 2002).

However, odor discrimination is also correlated with more parsimonious encoding models such as odor dependent patterns of spike bursting (Stopfer et al., 2003; Daly et al., 2004b; Lehmkuhle et al., 2006), which optimize on a relatively shorter time scale (ca. ∼140–240 ms after stimulus onset). This rapid timescale is consistent with correspondingly rapid and accurate behavioral responses in discrimination tasks in both insects and mammals (Uchida and Mainen, 2003; Budick and Dickinson, 2006; Wesson et al., 2008a,b). Spatial mapping of odorant responses in the OB and AL suggest that fine odor discrimination can be largely accounted for by subtle differences in glomerular input patterns alone (Johnson et al., 1998, 1999, 2005; Linster et al., 2001; Galán et al., 2004). Consistent with the notion that both spatial and temporal processes contribute to odor discrimination, we have recently shown that both odor dependent spatial patterns of glomerular output and spatio-temporal sequences of activations of these glomeruli can be observed during a response (Staudacher et al., 2009). In this case temporal patterns evolved over a ∼120-ms response window in an odor dependent manner. However, when data are collapsed across time, the spatial pattern alone remains odor dependent. Thus even within a given species, there are odor dependent spatial and temporal components that correlate to an animal's ability to discriminate.

The goal of the current study therefore was to better characterize oscillatory dynamics within the AL of the moth *Manduca* and to specifically establish whether odor-driven oscillations drive spike timing in a manner consistent with the transient oscillatory model. Specifically, we wanted to test the following predictions of this model: (1) Odor-driven LFPOs should be distributed across the AL; (2) The frequency content of these oscillations are: (a) of a consistent and stable frequency; (b) independent of the odor delivered and (c) occur reasonably consistently in response time across stimulus repeats; (3) unitary spiking responses should phase lock to oscillations; (4) this phase locking should be enhanced during odor-driven responses, relative to spontaneous activity; and (5) the local inhibitory network should establish oscillatory activity in the field and hence affect precise spike timing. To test these hypotheses, we placed 16 channel silicon microelectrode arrays within the AL to record spatial and temporal patterns of spiking activity across ensembles of sorted neural units. In addition, we recorded local field potentials (LFP) in multiple locations across the electrode array to investigate whether responses were localized or distributed. These recordings were made during both spontaneous activity and in response to odor stimulation in order to compare oscillatory and spiking activity during spontaneous epochs (i.e., ongoing processes of an active system), and during odor-driven responses. This allowed us to test whether the evolution of odor-driven oscillatory activity was associated with increased phase locking of individual cells. Finally, in order to address the functional role of local inhibition, the above observations were made before and during $GABA_A$ receptor blockade.

MATERIALS AND METHODS

MOTHS

Male *Manduca* moths were reared at West Virginia University using standard rearing procedures (Bell and Joachim, 1976). At stage 18 of pupal development, pupae were placed individually into paper bags and stored in a Percival incubator at 25°C with a reversed 16:8 LD cycle and 75% relative humidity. Moths were allowed to develop 5–7 days post eclosion, prior to experimental use; this ensures complete development of the AL and is consistent with our other behavioral studies.

NEUROPHYSIOLOGICAL PREPARATION

The method used for preparing moths has been described in detail elsewhere (Daly et al., 2004a,b). Briefly, individual moths were placed into a 14-mm ID copper tube and their heads firmly fixed to this tube with molten soft dental wax. An approximately 2 mm^2 window was cut into the top center of the head capsule. Next, cuts were made around the perimeter of the bilateral pair of pharyngeal dilator muscles (again ∼2 mm^2). The pharyngeal dilator muscles with attached head capsule cuticle were then slid forward into the previously cut window. The repositioned section of cuticle was then glued into position with superglue. By moving these feeding muscles, the brain can be directly accessed while leaving the animal functionally completely intact (Daly et al., 2004a). The prepared moth was then positioned on an air table and standard physiological saline was applied to the opened head to remove hemolymph and keep the brain moist. The flagellum of the ipsilateral antenna was then placed into a glass sleeve (ID 2.5 mm), which was connected to an olfactometer. Finally, the 16 channel electrode array (Neuronexus 2 × 2 tet) was placed into the AL using a high resolution motorized micromanipulator and controller (WPI HS6-3) under visual control. The probes were placed along a visually identifiable anatomical boundary between the AL and the rest of the brain, so that the probe array was located at the very caudal boundary of the AL. The depth of the deepest electrode in the array was between 500 and 600 μm, although this varied somewhat from moth to moth and was optimized for each recording. The completed preparation provided stable recordings for several hours. All recordings were performed between 1 and 4 h after the start of the moth's subjective night when they are normally active.

ODOR AND DRUG DELIVERY

Three microliters of monomolecular odorants were individually placed, undiluted, onto a ∼3-mm × 30-mm strip of Whatman filter paper and placed into a ∼1.7-ml glass cartridge (see **Table 1** for odor list). The cartridge ends were made of 1/16 in ID nylon

Table 1 | List of odorants used.

Odorant	Source	Purity	Density
1-pentanone	Sigma	99%	0.815
1-hexanol	Sigma	97%	0.861
1-heptanol	Sigma	98%	0.822
1-octanol	Sigma	99%	0.82
1-nonanol	Fluka	98%	0.88
1-decanol	Sigma	97%	0.89
2-hexanone	Sigma	98%	0.81
2-octanone	Sigma	98%	0.82
2-nonanone	Sigma	99%	0.821
2-decanone	Sigma	98%	0.824

luer fittings. Prepared cartridges were then connected to the odor delivery system such that the cartridge was only 2 cm from the tip of the antenna. Under normal conditions, dried and charcoal filtered air passed through the air supply tubing to the antenna at a velocity of 280 cm/s. This velocity was more than sufficient to minimize boundary layer effects around the antenna's sensilla, yet well within the normal flight speed of this moth (Stevenson et al., 1995). To stimulate with odor, clean air was shunted from the normally open line, which blows constantly on the antenna, to a second line which contained the odor cartridge, via a three-way valve (The Lee Co., model LFAA1200118H). The clean air and odor lines then merged into a T. The third arm of the T connected to the glass sleeve, in which the antennal flagellum had previously been inserted. Thus, air from either the clean air line or the odor line was continuously flowing over the flagellum.

Nine monomolecular odors were presented separately in random order. Each odorant was presented in 20 consecutive 100 ms pulses, each pulse spaced by 10 s. One hundred millisecond pulses were used based on empirical evidence indicating that this is the approximate mean exposure time for moths in the natural outdoor environment (Murlis and Jones, 1981) and a 10-s inter stimulus interval is sufficient to eliminate any modulation of responses due to repeated exposures in this model system (Daly et al., 2004b; Staudacher et al., 2009). This presentation scheme was repeated prior to, during and following bath application of 200 μM bicuculline methiodide (BMI) in physiological saline solution. This dose level has been shown to be effective in this (Waldrop et al., 1987; Christensen et al., 1998) and other insect species (Stopfer et al., 1997; Hosler et al., 2000). Drug treatment was initiated 10 min prior to the start of the second block of odor stimulations.

In a subset of two recordings, we also varied stimulus duration (50, 100, 500, and 1000 ms) to assess the effect of both briefer and more prolonged stimulation on the patterns of oscillatory activity. In this case we presented a restricted number of odors and did not apply $GABA_A$ receptor antagonists. We again presented 20 repeats of each stimulus/duration; each stimulus was again spaced by 10 s.

PHYSIOLOGICAL RECORDING AND SPIKE SORTING

Data from the four tetrodes on the electrode array were recorded using a 24 channel Neuralynx amplifier array and Cheetah data acquisition and stimulus control software. Spike data were recorded at 32 kHz using the tetrode sampling technique and local field recordings were sampled at 10.7 kHz from one electrode in each of the four tetrodes. **Figure 1A** highlights the approximate size relationship of the AL with the microelectrode array and identifies the relative position of the LFP recording sites on the array. In this case, LFP recordings were made from the top and bottom-most electrode sites to provide the greatest vertical coverage (250 μm deep by 200 μm wide); note that this represents approximately one-half the width of the AL. Spike channels were amplified between 2000 and 10,000 times, whereas all LFP recordings were amplified 2000 times. Hardware band pass filter settings were 600–6000 Hz for spike data and 1–125 Hz for LFP data.

Spikes were sorted in Matlab using the BubbleClust toolbox supplied by Neuralynx. This tool box implements a *k*th nearest neighbor algorithm in a 12 dimensional spike waveform feature space. BubbleClust identified clusters of spikes that were subsequently assessed and "cleaned" manually based on a number of descriptive and parametric statistics used to assess the overall quality and statistical separation of each cluster (see Daly et al., 2004b for complete description). Only those clusters that were significantly distinct from all other clusters and stable for the duration of the recording were considered for further analysis.

Local field potential recordings were re-filtered off-line via a windowed-sync FIR band pass filter provided by Neuralynx, using a rectangular smoothing window and 800 taps. This filtering procedure caused a linear phase lag that was then corrected to ensure accurate alignment of the LFPs with the spike data. Results were compared to several other filter implementations (including Butterworth and Chebyshev filters) to confirm that any oscillatory activity was not attributable to filtering artifacts, such as ringing. All implementations produced essentially identical results. Finally, for time frequency analysis (below) we also implemented a 60-Hz (±2 Hz, unless otherwise noted) windowed-sinc notch filter to the LFP data to remove line noise.

ANALYSIS

Several analytic techniques were used in order to quantify the relationship between parallel recorded LFPOs and between LFPOs and unitary spiking. These methods include LFP–LFP and unit–LFP cross correlations, vector strength analysis, power spectral density (PSD), cross spectral density (CSD), and time frequency response spectrograms (TFRs). In all cases a 1-ms binning window was used unless otherwise noted and all calculations were performed in either MATLAB or Neuroexplorer. Statistical analyses of the data generated from these methods were performed in MATLAB and SAS.

First, in order to characterize the frequency content of odor-driven oscillations, TFR spectrograms were generated via the Short Time Fourier Transform method. TFRs were calculated with the "tfrsp" function in the Time Frequency Toolbox for Matlab (freely available from the Centre National de la Recherche Scientifique at http://tftb.non-gnu.org). This analysis was performed on individual responses to odor stimulation and the results averaged to highlight consistent oscillations whilst averaging out inconsistent oscillatory activity. Because TFRs were calculated individually for each response, then averaged across repeats, the approach abstracts from (and is hence tolerant of) variation in oscillatory phase between stimulus-aligned traces. This is important because averaging of raw oscillations from repeated trials for example could potentially cancel out oscillations that are not phase aligned across trials. TFR results were *z*-score normalized to highlight possible differences in frequency content and relative amplitude across the four LFPO recording sites (see **Figure 1A**). Normalization also allowed comparison of pre- versus post- $GABA_A$ receptor blockade in order to highlight changes in frequency content for these specific comparisons.

To assess whether each odorant produced a unique frequency modulating (FM) sweep pattern (i.e., differences in frequency range, onset/offset, and duration) we implemented a discriminant analysis using a supervised classification approach. Specifically we used the support vector classifiers method (Boser et al., 1992; Galán et al., 2004; see also Methods in Appendix). Here the goal was to determine if TFR's generated from individual odor responses could

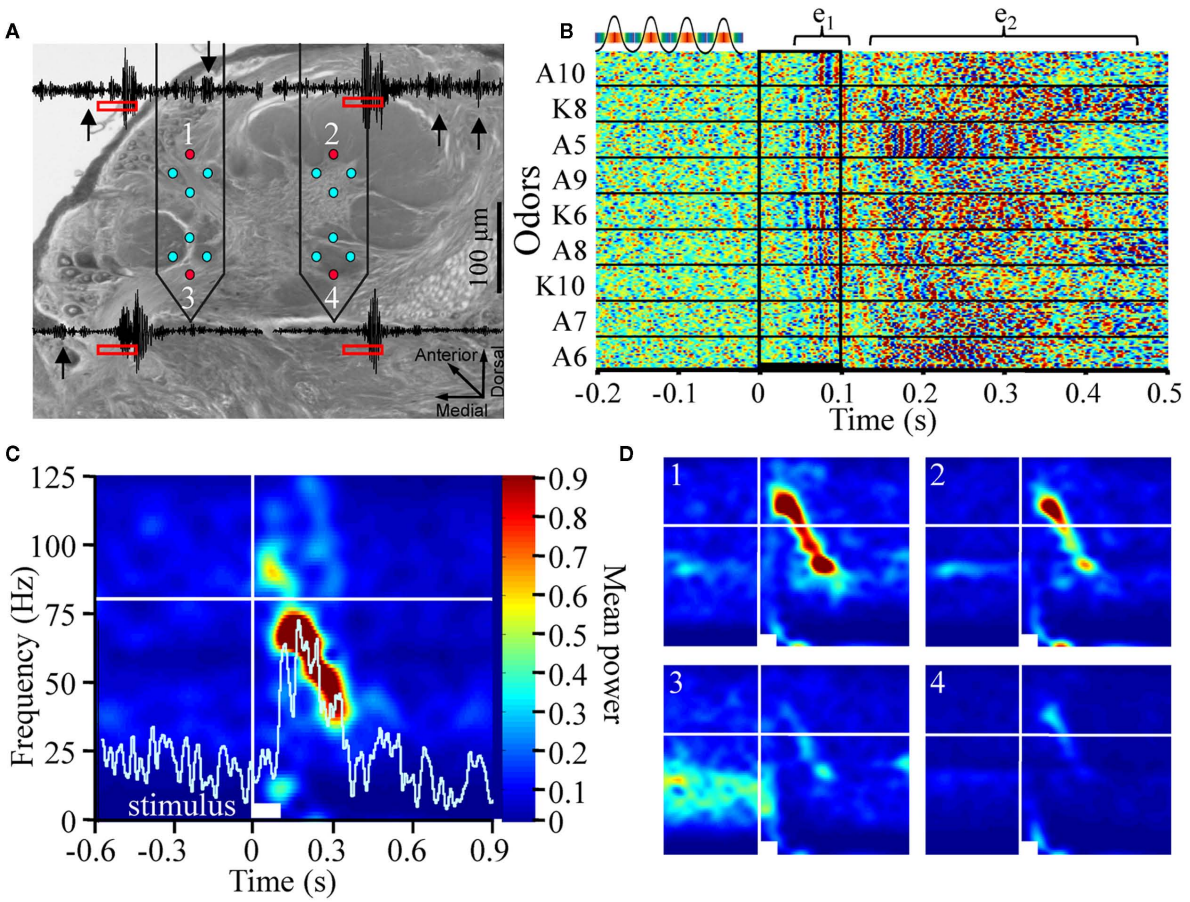

FIGURE 1 | Odor-driven oscillations are frequency modulating and localized to a subset of recording sites (A). Schematic of the multichannel electrode array configuration within the antennal lobe (AL; Cason stain). Neuronexus two-shank 2 × 2 tetrode silicon electrodes array, each shank containing eight electrodes, which are arrayed into two tetrodes (for tetrode sampling of spiking activity), are placed into the AL. One electrode from each tetrode was used to record local field potentials (LFP; red circles). This provided four parallel LFP recordings (inset) from across the AL. Note in these example traces, which have been band pass filtered to highlight 25–100 Hz that there are several instances of spontaneously occurring "bouts" of oscillatory activity (inset black arrows). **(B)** Rasterized peri-stimulus LFP traces from 20 presentations of 9 different alcohols and ketones. Each trace has been converted to a one dimensional color coded line to indicate whether the voltage was high (reds) or low (blues; see above inset). All 180 responses are aligned by stimulus onset (stimulus duration indicated by vertical black lines) and stacked to create a single composite panel. Odors are separated by black horizontal lines and identified (left of panel) by moiety (A = alcohol; K = ketone) and carbon chain length (5–10). Highlighted above are two distinct oscillatory response epochs; an early phase (e_1), and a later phase (e_2). **(C)** Averaged time frequency response spectrogram showing consistent oscillatory power (normalized color scale) as a function of frequency (*Y*-axis) and peri-stimulus time (*X*-axis) for a single odor (1-pentanol). Displayed are the averaged results of 20 independent TFR analyses for each of the 20 presentations 1-pentanol. Inset white bar (bottom) indicates stimulus duration. Vertical white line indicates odor onset. Horizontal white line is an 80-Hz frequency reference. Inset histogram (light blue) is the corresponding mean population spiking response to highlight the relationship between the onset and duration of the spiking response and the oscillations. **(D)**. Comparison of the spectrogram results from the four parallel recording sites in response to 2-octanone. Inset numbers correspond to electrode recording site shown in **(A)**. Power is normalized to the same range to make panels directly comparable. All panels are scaled and referenced as in **(C)**.

be classified accurately into odor groups based on the multidimensional shape of the FM sweep. Briefly, we first implemented a data reduction step. This involved cropping the TFRs generated from each individual stimulus to a 12 to 124-Hz frequency range, and a −30 to 1000-ms time range and then defined a region of interest (ROI) within each TFR, consisting of all pixels that were above the 99th percentile of power density. Next we generated a filtering mask which represented all pixels from all odors that were above the 99th percentile and applied it to all TFRs. This data reduction step effectively selects all significant frequency–time content both within odor repeats and between different odorants and removes all irrelevant data. We then implemented the support vector classifiers using both a linear and third order polynomial kernel functions. This analysis was performed on two recordings. Finally, we calculated an experiment wide 95% confidence interval to identify classification performance as significantly above chance.

In order to statistically quantify differences in odor-driven spectral content prior to and during $GABA_A$ receptor blockade, PSD analysis was performed. For this analysis we calculated the

PSD for each odor-driven response in a 1000-ms window starting at stimulus onset. Note that in response to a 100-ms stimulus, TFR analysis indicated that most evidence of odor-driven oscillatory activity terminated within ∼500 ms, thus all odor-driven oscillatory activity should be encapsulated within this 1000 ms window. Furthermore, TFR analysis indicated that odor-driven responses were typically no higher than 100 Hz, thus PSD analysis was performed on a 1 to 100-Hz frequency range, in 0.5 Hz steps using Neuroexplorer. Resulting PSD data were then statistically analyzed using ANOVA in SAS to determine if spectral power at each frequency step was affected by $GABA_A$ receptor blockade.

To quantify the precise relationship between LFPs recorded from different sites across the AL, we calculated all possible pair wise LFP–LFP cross correlations. For this analysis, a ±100-ms sliding window was used and data was sampled from spontaneous and odor-driven activity (−1000 to 0 and 0 to 1000 ms respectively; time relative to stimulus onset). This analysis was performed individually for each stimulus presentation, both before and during $GABA_A$ receptor blockade, for all recordings. The cross correlation not only indicates whether any two traces are correlated, but also highlights any temporal structure of that correlated activity. In order to extract and quantify the temporal structure embedded in the cross correlation, CSDs were calculated for each individually calculated cross correlation. CSDs specifically measure the relative power of any periodic structure present in the cross correlation, hence the CSD quantifies oscillatory activity that is distributed across the two recording sites used for a given cross correlation. Results of the CSD analysis were then statistically analyzed using ANOVA in SAS to determine if distributed power at each frequency step varied as a function of spontaneous versus odor-driven sampling time, and as a function of $GABA_A$ receptor blockade.

We used two related methods to quantify the phase relationship between unitary spiking patterns and the LFPO, as both a function of spontaneous and odor-driven responses, and as a function of $GABA_A$ receptor disruption. First, the instantaneous phase of the LFP was calculated in MATLAB as the analytical phase of the signal, i.e., as the instantaneous angle of the complex time series obtained from the Hilbert transform of the LFP. Results of this analysis provided a visual display, as a series of histograms for each cell, highlighting the distribution of spikes relative to oscillatory phase. Our second approach was to use vector strength analysis. Vector strength quantifies the degree to which unitary spiking activity exhibits a tendency to occur at a particular phase of the LFPO by calculating the reciprocal of the circular variance of the distribution of phases of sampled spikes to the LFP; this was also performed in MATLAB. For this analysis data was sampled from three peri-stimulus epochs; spontaneous activity prior to stimulation (−300 to 0 ms), an early response epoch (50–350 ms), and a late response epoch (350–650 ms). Note that 0–50 ms was not considered in this analysis to account for the time lag between when the odor valve opened and the earliest physiological evidence of AL responses. Furthermore, because preliminary results indicated that odor-driven oscillatory activity was relatively broad band and in a different frequency range than what was observed during spontaneous activity, we performed the vector strength analysis on LFP data that was band pass filtered to highlight two frequency domains; low (25–55 Hz) and high (55–85 Hz). Filtering LFP oscillations into narrower pass bands results in a less complex wave and hence allows more accurate calculation of spike phase. Filtering also allows segregation and characterization of phase locking in the separate frequency domains. Subsequently, ANOVA was used to determine whether vector strength was dependent on: peri-stimulus time (i.e., pre, early, and late response epochs), the odorant used, and, $GABA_A$ receptor blockade. Analysis was performed after accounting for random effects such as individual recording (moth) and individual differences between spiking units. Finally, the filter implementation (low versus high frequency pass band) was also included in the statistical model to establish whether the different pass bands produced statistically similar results.

Finally, to establish whether vector strength values were above what might be expected by chance, these values were recalculated for one of the recordings using a bootstrapping method, whereby spike times were first jittered randomly within a series of time ranges from ±1 to ±10 ms in 1 ms steps (spikes were re-jittered 30 times for each step in the jitter width). Results were statistically analyzed using ANOVA in SAS to determine if vector strength was dependent on the amount of jittering for the three peri-stimulus epochs.

RESULTS

ODOR-DRIVEN OSCILLATIONS ARE FREQUENCY MODULATING AND TYPICALLY LOCALIZED TO A SUBSET OF RECORDING SITES

The transient oscillatory model of odor identity encoding posits that odor stimulation should produce a reasonably narrow (hence stable) frequency band of oscillatory activity, which is distributed across much if not all of the AL. This prediction is based on the morphological characteristics of LNs, which arborize broadly and mediate oscillatory activity. The model furthermore predicts that the elicited frequency should be independent of the odorant presented (Laurent and Davidowitz, 1994; Laurent et al., 1996a). A stable response frequency that is independent of varying stimuli is an important aspect of this model because the second-order olfactory center, the MB, are proposed to integrate inputs from the AL across a limited time window defined by delayed feed-forward inhibition from the lateral horn (Laurent, 2002; Perez-Orive et al., 2002). Therefore, we detailed the precise nature of the oscillatory activity in the AL of *Manduca* to determine if these criteria could be met.

Figure 1A is a schematic depicting the approximate positioning and size of the electrode array in the AL. This figure also highlights which electrodes from the 16 channel array were used for LFP recordings. In this case, we made parallel LFP recordings from the highest two and lowest two recording sites (shown in red), during periods of spontaneous activity and in response to odor stimulation. Superimposed on **Figure 1A** are peri-stimulus LFP traces representing a range of typical 25–100 Hz oscillations. These oscillations occur both spontaneously (black arrows) and in response to a 100-ms odor presentation (stimulus highlighted by red rectangles). Note that oscillations can be clearly observed during spontaneous activity on three of the four traces, albeit at a lesser amplitude than those observed during an odor-driven

response. This suggests that oscillations are an ongoing product of a spontaneously active system.

Figure 1B displays rasterized peri-stimulus LFP traces from a single LFP electrode in response to the 20 repeats of each of the nine odors used, thus comprising a total of 180 rows of LFP data. First, in the 200-ms prior to odor stimulation, there is again evidence of spontaneous oscillations as indicated by the random "peppering" of interleaved red and blue "blips" in the horizontal (i.e., time) axis. The presence of spontaneous oscillatory field activity should not be particularly surprising given the spontaneously active nature of the insect AL (Perez-Orive et al., 2002; Stopfer et al., 2003; Daly et al., 2004b).

However, in the presence of odor, LFPOs become both larger in amplitude and in some cases coherent across repeats when aligned by the onset of odor stimulation. For example, inter-trial coherence can be seen in **Figure 1B** by the prominent vertical striping of blue and red across the 20 repeats of each odor as well as across odors in some instances. This indicates that consecutive oscillation peaks (red) and valleys (blue) were phase aligned from response to response. This "inter-trial coherence" across stimulus repeats is most prominent for 1-pentanol (A5) between 150 and 300 ms post stimulus onset. **Figure 1B** also suggests that odor-driven oscillations have two distinct epochs; these epochs are highlighted by inset brackets labeled e_1 and e_2. Early phase high frequency oscillations (e_1) emerge reasonably consistently (∼60 ms post stimulus onset) and last for ∼50 ms. Early phase LFPO's appear to some degree for most odors. This early phase oscillation precedes onset of the odor-driven I_1, a bicuculline sensitive, fast inward Cl^- conductance observed in *Manduca* PNs (Waldrop et al., 1987; Christensen et al., 1998). The I_1 phase occurs prior to spiking responses of PNs. Thus, based on the timing of early phase oscillations, they likely relate to sensory cell input and early synaptic processing that occurs prior to the onset of PN spiking responses.

Starting no earlier than ∼120 ms is a second bout of high frequency LFPOs (e_2), which appear less coherent in response time, though exceptions exist (again see responses to 1-pentanol). These later phase oscillations tend to be odor dependent. However, it is not always apparent, using this visualization method, whether odor-driven e_2 oscillatory activity is consistent from trial to trial in terms of onset, frequency, and/or phase.

To establish whether odor-driven oscillations are consistent in response time and produce a reasonably stable frequency, thereby providing a faithful temporal encoding mechanism that is independent of odor identity (Laurent and Davidowitz, 1994; Laurent and Naraghi, 1994), we calculated TFRs. **Figure 1C** displays the averaged results of 20 separate TFR analyses performed on each of the 20 presentations of 1-pentanol; results are based on a single recording site. Power was normalized to a probability distribution and color coded. The first and most striking observation is that odor stimulation drives an oscillatory response that is FM. In response to all 9 odors for all 10 moths in the study (1,800 total odor presentations) we observed that odor-driven oscillatory responses produced a high-to-low frequency sweep that typically ranged from 80 to 100 Hz peak frequency down to 20–40 Hz. Superimposed on this TFR is the corresponding peristimulus histogram, highlighting the summed spiking response for all units in the corresponding population. Results were then averaged across the 20 stimulus repeats. Consistent with the above observation that oscillatory activity precedes spiking activity in AL neurons, **Figure 1C** shows that high frequency oscillatory activity precedes the population spiking response by ∼60 ms in this case.

Across the 10 moths used for this analysis, odor-driven FM oscillatory responses were also spatially localized. **Figure 1D** displays typical TFRs generated from the four parallel LFP recording sites from a single animal in response to 20 presentations of 2-octanone. In **Figure 1D**, the upper left electrode (1) recorded an FM sweep starting at ∼60 ms post stimulus onset and a peak frequency of 100 Hz. The frequency of this oscillation then modulated down to 35 Hz by 400 ms post stimulus onset. This pattern is evident in the upper right electrode (2) as well, although the relative power was lower. In the lower two electrodes (3 and 4) however, almost no evidence of this FM oscillation was present. As highlighted in **Figure 1D** electrode 3, there were also consistent spontaneous 20–50 Hz oscillations that were interrupted upon odor stimulation; this form of spontaneous oscillatory activity was observed in at least one of the LFP recording sites for 6 of the 10 moths that we analyzed using the TRF method (see **Figure A1** in Appendix for examples from all six recordings). As suggested by the raw traces in **Figure 1A**, these ongoing oscillations tended to be spontaneous but intermittent; this accounts for the patchy power structure observed in **Figure 1D**. Thus, while 2.7 ± 1.15 (mean ± SD) recording sites exhibiting oscillatory activity of some sort, only 1.8 ± 1.2 (mean ± SD) recording sites actually displayed the same odor-driven FM response patterns as seen in electrodes 1 and 2. Contrary to the predictions made of the transient oscillatory model, this suggests that odor-evoked oscillations that were FM and typically constrained to a sub region of the AL.

ODOR-DRIVEN OSCILLATIONS MODULATE IN A STIMULUS SPECIFIC MANNER

As mentioned above, the transient oscillatory model posits that the oscillations should be frequency invariant across different odors (Laurent and Davidowitz, 1994; Laurent et al., 1996a). TFR results, however, suggest that odor-driven FM LFPOs were dependent on both the physical features of the odorant molecules used as well as the duration that a test odorant was presented. **Figure 2A** displays TFRs from a single recording electrode in response to each of the nine odors used (see also **Figure A2A** in Appendix). Responses driven by both the short chain alcohols and ketones exhibited a maximum frequency of approximately 100–105 Hz. As the carbon chain length of the stimulus was increased the peak frequency at response onset appeared to drop. For example, **Figure 2** shows a maximum frequency of approximately 85–95 Hz when the moth was stimulated with 1-decanol and 2-decanone respectively. The duration of the oscillatory response also increased with carbon chain length. Note for example, the longer chain odors modulate more slowly, in the case of K10 the response persists across the time window whereas the shorter chain odors do not. Finally, this implies that the rate of the FM sweep (that is, how long it took to modulate from high-to-low frequency), is slower with the

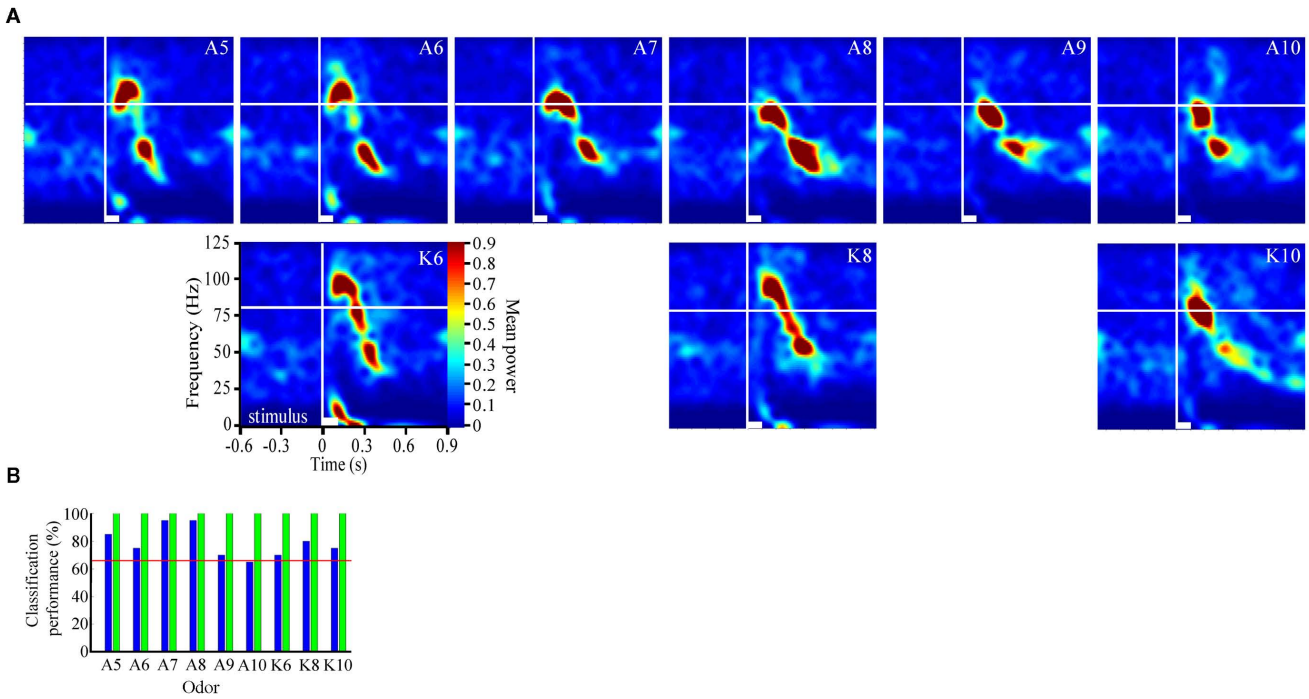

FIGURE 2 | Odor-driven oscillations modulate in an odor-specific manner. **(A)** Time frequency response spectrograms calculated as a function of changing molecular features of the stimulus. All panels are the result of TFR analyses that were calculated, normalized, and presented as described in **Figure 1**. Each panel represents the averaged TFR result for a single odorant (inset and identified as described in **Figure 1**). Note that alcohols and ketones of common length are aligned into columns. Note that with the longer carbon side chains, the initial peak frequency is decreased by ∼15–20 Hz and the FM sweep appears both slower and longer in duration (see also **Figure A2** in Appendix). **(B)** Results of the discriminant analysis, showing classification performance for each odorant as a function of both the linear (blue bars) and polynomial discriminant functions (green bars). Red line highlights the 95% confidence limit and hence identifies those odorants for which the linear and third order polynomial kernel functions were able to accurately classify significantly better than chance. Note that while the linear kernel in this case was not able to classify A10 significantly better than chance, the polynomial kernel could discriminate all odors with 100% accuracy (see also **Figure A2** in Appendix).

longer carbon chain odors. This resulted in a shallower and more prolonged downward sweep.

To assess the odor dependency of FM responses we performed a discriminant analysis using support vector classifiers and implementing both a linear and polynomial discriminant (or kernel) function on a subset of two recordings. The results of this analysis for the dataset shown in **Figure 2A** are displayed in **Figure 2B** (for the second example see **Figure A2** in Appendix). The linear discriminant function (blue bars) was able to accurately discriminate all odors significantly above chance (red line indicates 95% confidence level) for all but a single odor (A10), whereas the polynomial function (green bars) effectively classified with 100 percent accuracy for all odors. This indicates that each odorant produced a unique frequency response "signature." This odor dependency can be interpreted as either volatility-dependent or carbon chain length-dependent. In either case, these results indicate that the qualitative nature of the oscillation is clearly stimulus-dependent.

In two moths we also varied stimulus duration for a subset of odors. **Figure 3** displays TFRs results from an individual recording site where the odor 2-hexanone was presented 20 times for each of four different stimulus durations, ranging from 50 to 1000 ms. Note that there is a stimulus duration dependent elongation of the FM sweep; this was observed in both moths where duration was manipulated. Thus again, in contrast to theoretical expectations, the frequency content, and rate of modulation are clearly affected by the odorant used as well as the duration of the stimulus.

CORRELATED FIELD ACTIVITY IS STATE DEPENDENT AND ONLY WEAKLY PERIODIC

The observation that odor-driven FM oscillations are localized to a subset of recording sites does not necessarily imply the absence of more subtle distributed oscillatory activity that the TFR method was unable to quantify. Another method for identifying distributed activity across the four LFP recording sites was to directly compare pairs of sites using cross correlation analysis. Results of this analysis revealed that while different recording sites are in fact moderately to highly correlated, this correlation decays rapidly in time, producing very little evidence of periodic structure.

For example **Figure 4A** displays the cross correlation between LFP recording site 1 with all other LFP recording sites for a single animal. This analysis was based on one 40 min recording segment containing both spontaneous and odor-driven activity and serves to highlight that recording sites are clearly correlated at 0 s. Indeed, inspection of these cross correlations reveals that the peak correlations are centered on 0 s. and range from $r = 0.33$ (sites 1 and

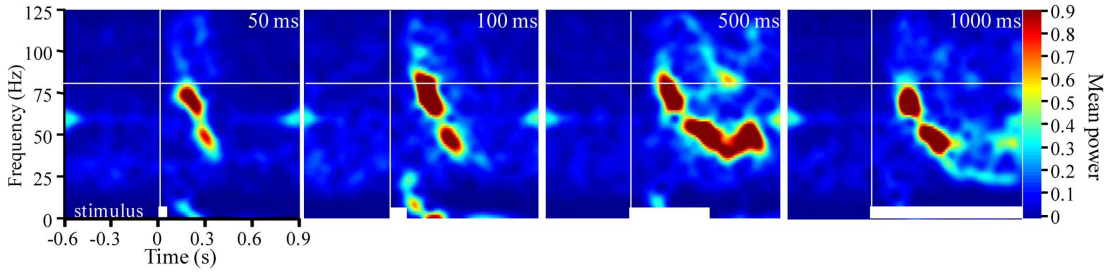

FIGURE 3 | Frequency modulation is dependent on stimulus duration. Stimulus averaged TFR spectrograms for a single odor (2-hexanone) presented for durations ranging from 50 to 1000 ms (20 presentations/duration). White bars at the bottom of each panel indicate stimulus duration. Vertical and horizontal white lines reference stimulus onset and 80 Hz respectively. Note that with increased stimulus duration there are more prolonged FM sweeps.

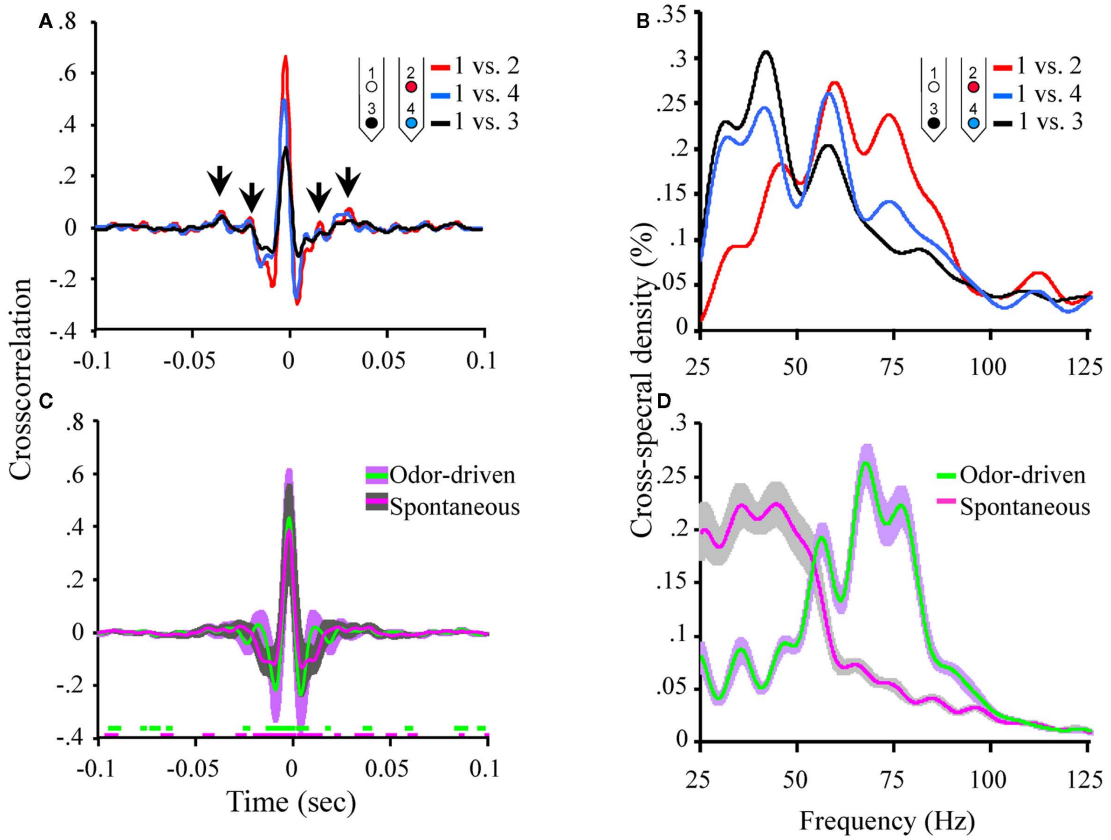

FIGURE 4 | Oscillatory local field activity is weakly distributed and contains state dependent frequency content. (A) The cross correlation between LFPs recorded from site 1 (tetrode 1) and the three other LFP recording sites in the tetrode array (see **Figure 1A**). Inset schematic of the 2 × 2 tetrode array shows the respective LFP recording site positions. For purposes of simplification, only the site on each tetrode that was used for LFP recordings is shown. Sites are color coded to identify which cross correlation trace is associated with which recording electrode. Thus for example, the red trace is the cross correlation between site 1 (white) and site 2 (red). Results based on a total of ∼20 min of continuous recording time from a single moth and contain both spontaneous and odor-driven data. Inset arrows identify the periodic structure observed in the cross correlation between electrodes 1 and 2. Note that these small "off center" peaks in the cross correlations occur in a periodic manner suggesting the presence of a weak but distributed oscillation. **(B)** The corresponding cross spectral densities for the data shown in **(A)**. These CSDs indicate that the cross correlation between electrodes 1 and 2 contained higher frequency content than 1 and 3 and 1 and 4 respectively. **(C)** Mean cross correlations for spontaneous versus odor-driven periods. Results based on all possible pair wise comparisons between the four LFP recording sites across seven moths. Odor-driven cross-correlation based on 1 s samples, starting at odor onset, for each of 180 odor stimuli per moth. Spontaneous cross correlation based on 180 1 s samples taken immediately prior to odor onset for each odor stimulus, per moth. Note that cross correlations were calculated for each stimulus and results were averaged. Shaded regions around the mean represent ±1 SD. Inset color coded broken lines above the *X*-axis indicate periods where a significant difference in the cross correlation value from 0 was observed. Significance threshold for these tests was set to 2.3×10^{-6} to maintain an overall alpha of 0.01 (see **Figure A3** in Appendix for an expanded view). **(D)** Mean cross spectral densities as a function of spontaneous and odor-driven activity. Results based on all responses from the same seven moths shown in **(C)**. Shaded area represents SE.

3), to $r = 0.68$ (sites 1 and 2). This high degree of correlation occurs between sites which are as much as 250 μm apart. As highlighted in **Figure 4A**, the magnitude of these correlations does not appear to be dependent on the distance between recording sites. For example, the cross correlation between sites 1 and 4 in this experiment (250 μm) was considerably higher than between sites 1 and 3 (200 μm).

Note too that there are some small but regular "ripples" in the individual cross correlations (highlighted by inset arrows in **Figure 4A**). These small and roughly evenly spaced (in time) peaks and valleys, indicate that there is indeed some periodic temporal structure in the cross correlation. However, the peaks of these correlations are quite small (ca. $r < 0.10$). Nevertheless, by calculating the CSD of each individual cross correlation, any periodic temporal structure in the cross correlation can be quantified; the corresponding CSDs for the cross correlations displayed in **Figure 4A** are shown in **Figure 4B**. **Figure 4B** plots the amount of distributed oscillatory power by frequency and indicates that there was correlated periodic activity between sites 1 and 2 (red trace) centered around ∼60 Hz but spanning broadly from ∼40 to 85 Hz. By contrast, distributed oscillatory power present in the cross correlations between sites 1 and 3, and sites 1 and 4 contained frequency content in a lower frequency range indicating that the correlated oscillatory activity between sites 1 and 2 is not the same as that found between sites 1 and 3 or 1 and 4.

Figure 4C summarizes the mean of all possible unique pair wise cross correlations for all recordings used in this analysis (six possible unique pair wise comparisons per moth and seven moths used for this analysis). Results are based on a total of 360 1000 ms samples and are broken down as a function of spontaneous versus odor-driven activity. Mean cross correlations have been plotted with ±1 SD shading to provide a sense of variability in the cross correlations; this indicates that the greatest variability occurs in a ±50-ms window centered around 0 ms.

The mean peak cross correlation was 0.38 (SD = ± 0.18) for spontaneous activity and 0.43 (SD = ± 0.19) for odor-driven activity. A one-tailed paired *t*-test comparing the peak cross correlation revealed that odor produces a significant increase in this measure across the recording sites ($p < 0.001$). However, because these cross correlations are presented as means of multiple analyses, all but the most consistent and robust temporal details tend to be averaged out. To determine whether the secondary peaks (and valleys) that flank the central peak at 0 ms were significantly different from $r = 0.0$, a two tailed, *t*-test was calculated, for each frequency step; this was performed for both the spontaneous and odor-driven cross correlations. Inset along the *X*-axis in **Figure 4C** are the results of this test. In this case breaks in the colored lines indicate moments where the spontaneous or odor-driven cross correlations were not statistically different from 0. Conversely, the presence of colored lines highlights where there were significant deviations from 0. Thus, in spite of the fact that the flanking peaks and valleys were quite small, there were periods where the correlations were significantly different from 0, suggesting that there is indeed some correlated temporal structure that is distributed across pairs of recording sites.

Therefore, in order to quantify this distributed temporal structure, CSD's were calculated for each individual cross correlation. Results of these CSD analyses were then averaged as a function of spontaneous versus odor-driven epochs and are displayed in **Figure 4D**. As shown in **Figure 4D**, the most striking difference between spontaneous and odor-driven CSDs is that they produce distributed temporal structure in nearly dichotomous frequency ranges. That is, during spontaneous epochs of time, the bulk of distributed periodic activity was in the range of ∼25–55 Hz. In the presence of odor however, there was relatively more distributed power in the range of ∼55–85 Hz. After accounting for all statistical main effects, ANOVA revealed distributed oscillatory power was significantly dependent on the interaction between the presence or absence of odor and frequency ($p < 0.0001$). Collectively then, this pattern of results indicate that both spontaneous and odor-driven activity produce weak but detectable distributed oscillatory activity in different frequency ranges. Our expectation was that the relationship between oscillations across recording sites would be far greater than actually observed; nevertheless the presence of distributed oscillations is consistent with the transient oscillatory model. The next question is whether these weak but distributed oscillations affect spike timing in a pattern consistent with the model.

UNITARY SPIKING PHASE LOCKS TO LFPOS AS AN ONGOING, NOT ODOR-DRIVEN PROCESS

The odor-driven increase in the amplitude and frequency of weakly but distributed oscillatory activity across the AL could provide enhanced synchrony of distributed spiking behavior from populations of AL neurons as predicted by the transient oscillatory model (Laurent and Davidowitz, 1994; Laurent et al., 1996a,b). This would require that phase locking of spikes to the oscillations occurs during odor-driven responses and this phase locking would presumably be stronger than what is observed during spontaneous activity. To test this hypothesis, we calculated unitary vector strength for the three peri-stimulus time epochs. These calculations were made for each unit across all stimulations of all odors. Thus for each moth, individual unitary vector strength calculations were based on 180 ms × 300 ms (or 54 s) samples for each of the three peri-stimulus time epochs. Furthermore, given that two relatively distinct frequency domains were observed for distributed oscillatory activity during spontaneous and odor-driven epochs (see **Figure 4D**), vector strength calculations were made based on both the low (25–55 Hz) and high (55–85 Hz) LFP band pass filter implementations. Finally, only vector strength values that were calculated based on 60 or more spikes were included in the analysis (376,518 spikes, or 87% of the 433,667 spikes sampled). This criterion was used because vector strength describes the variation in a distribution of spikes relative to the phase of a 360° oscillation cycle. Hence, with low spike counts, the vector strength measure is less reliable, particularly in cases where there are spikes that are outliers. To analyze variation in vector strength, we implemented a mixed General Linear Model in SAS using the GLM procedure to model data from five moths (90 neurons) for which we had comprehensive datasets (including both before and during $GABA_A$ receptor blockade via BMI application; BMI results described below).

The overall model was significant, explaining 49% of the variance in unitary vector strength scores ($p < 0.0001$). This analysis

found that the main effects of odor identity (of the nine odors used), the peri-stimulus time epoch (−300 to 0; 50–350; and 350–650 ms, relative to onset of odor stimulation), and drug treatment (before versus during BMI application) all significantly influenced vector strength values ($p < 0.0001$). However, the LFP frequency range used to calculate the vector strength values (25–55 or 55–85 Hz) was not significant ($p = 0.7274$) indicating that vector strength values were approximately the same for both frequency domains.

Of particular importance is the significant main effect of peri-stimulus epoch, which compared vector strength as a function of the three different epochs of peri-stimulus time. *Post hoc* analysis of this effect (**Figure 5A**) indicates that, independent of the band pass filter implementation, spontaneous activity prior to odor stimulation (−300 to 0 ms relative to odor onset) produced the greatest overall vector strength values. Interestingly, from 50 to 350 ms, the peri-stimulus epoch typically associated with odor dependent spiking patterns (Stopfer et al., 2003; Daly et al., 2004b; Staudacher et al., 2009), produced about half the vector strength, indicating relatively weak phase locking during odor-driven responses. Even during the late response epoch (350–650 ms), while there was a significant increase in vector strength relative to the early response epoch ($p < 0.01$), it was still significantly lower than vector strength during spontaneous activity ($p < 0.01$). In order to better understand this "response epoch" effect at the single unit level, phase lag histograms are displayed for all 26 units from a single recording and in response to 20 presentations of 1-decanol, (**Figure 5B**). Results are displayed for the same three peri-stimulus epochs. Here we show that ongoing (i.e., spontaneous) spiking tends to occur, in this recording, on the rising phase of the oscillation cycle. In this exemplar case, all units show a decrease in their phase relationship to the oscillation during the

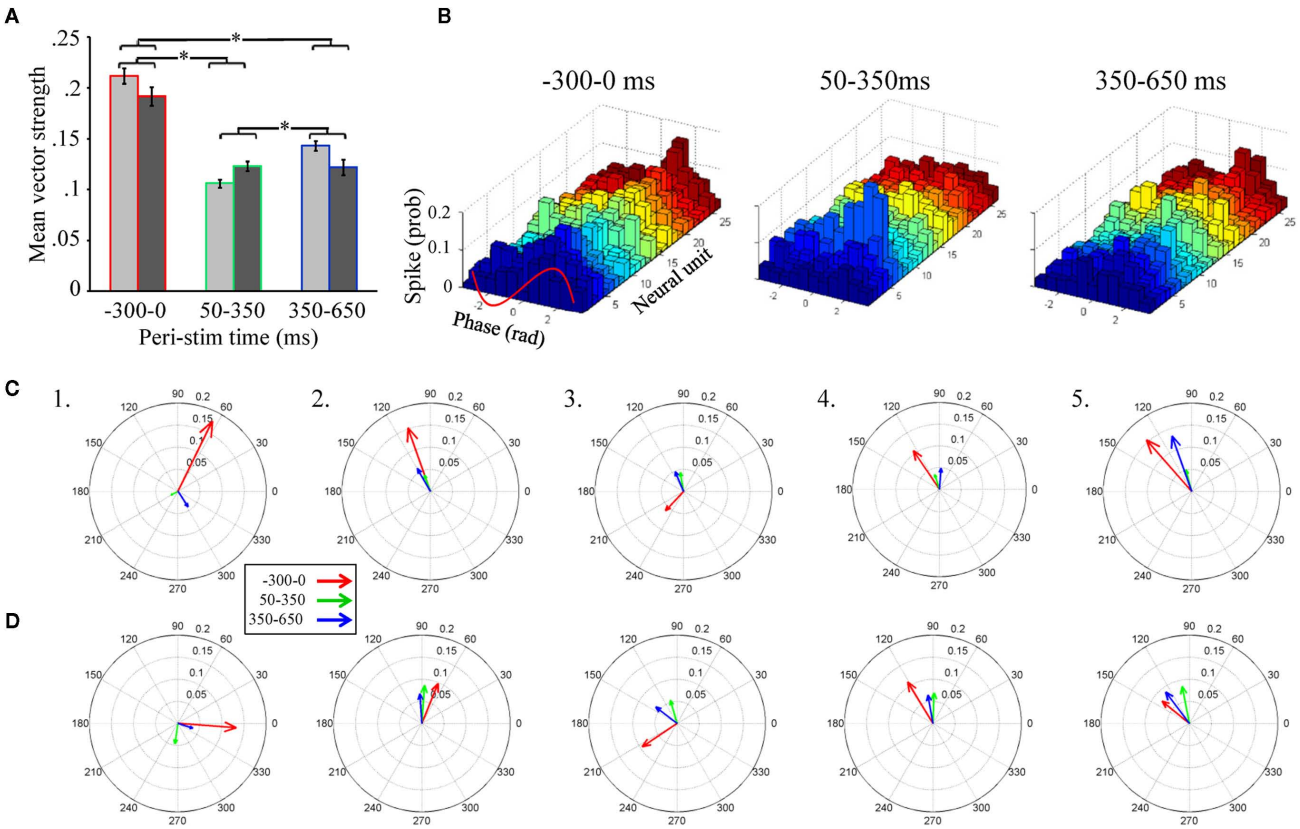

FIGURE 5 | Unitary spiking phase locks to LFPOs as an ongoing not odor-driven process (A). Mean unitary vector strength as a function of the three stimulus epochs for the five moths used in this analysis. Results are broken down by LFP band pass filter range. However, note that there are no significant differences within any peri-stimulus response epoch between the 25–55 Hz (light gray) and 55–85 Hz (dark gray) band pass filter implementations. Inset are the results of statistical *post hoc* comparison mean vector strength values (averaged across filter implementation) across the different peri-stimulus time epochs; means with significant differences are indicated with an asterisk. Errors bars represent SE. **(B)** Phase lag histograms for all 26 recorded units from a single recording. Results based on one of the four LFPs recording sites, and in response to the 20 presentations of 1-decanol. Results are broken down by spontaneous (300 ms just prior to odor onset), early response (50–350 ms), and late response (350–650 ms). Note that several units (but not all, for example see the light green histogram) show a tendency to spike on the rising phase of the oscillation cycle (between 0 and 1.57 radians; see also inset red wave) during spontaneous activity. During odor-driven responses only one unit retains a strong (increased) phase relationship to the oscillation while the bulk of the units lose their phase relationships. **(C)** Population-level vector strength as a function of the three stimulus epochs. Results are for the 25 to 55-Hz band pass filter implementation and each panel represents a different animal. **(D)** The same as **(C)** except using LFP data that was 55–85 Hz band pass filtered.

early phase of the odor-driven response except one (**Figure 5B**, 50–350 ms; unit 7, light blue histogram), which increased and shifted phase essentially to the peak of the oscillation cycle. Finally, during the late phase of the response, units appear to come back into a similar phase alignment as was observed during spontaneous activity.

In summary, of those units meeting the 60 spike minimum criterion for both spontaneous and early response epoch ($N = 65$ units), odor stimulation resulted in 68% of those units producing a 0.09 ± 0.08 (mean $\pm$ SD) reduction in vector strength relative to spontaneous activity, while only 32% produced a 0.03 ± 0.04 (mean $\pm$ SD) increase. Furthermore, consistent with the finding that vector strength recovers during the late response epoch, we found that 71% of units produced a 0.07 ± 0.05 (mean $\pm$ SD; $N = 68$ units) increase in vector strength scores from the early to late response epoch, while only 29% of cells produced a 0.06 ± 0.08 (mean $\pm$ SD; $N = 68$ units) decrease. Thus, more than two-thirds of units show a relative decrease in vector strength during the early response epoch and recover during the late response epoch.

Unitary vector strength measures represent the best possible scenario for highlighting the relationship between LFPO's and spike timing of individual units because it is tolerant of differences in preferred phase angles between units. However, the transient oscillatory model predicts that not only do spikes phase lock to oscillations during an odor-driven response, but that phase locking of the responding assembly of output cells results in an increase of synchronized spiking of the population on an oscillatory timescale. In order to understand how oscillations might synchronize a population of parallel recorded neural units, the same vector strength analysis can be applied to a population vector, which is the summed spiking behavior of the recorded population. This population vector strength approach is sensitive to differences in preferred phase angles between units. Thus, as the variation in preferred phase angle among individual neurons in the population increases, the population-level measure of vector strength (and hence synchrony among individuals) will decrease. **Figures 5C,D** display population-level vector strength and preferred phase angle across all recorded units for the five recordings used in the statistical analysis. We again used the same 60 spike minimum for inclusion into the population vector. Each panel in **Figures 5C,D** displays the vector strength for the same three peristimulus response epochs for a single moth; results are collapsed across all presentations of all odors. **Figure 5C** displays results based on the 25–55 Hz band pass filter implementation, whereas **Figure 5D** is based on the 55–85 Hz filter. As statistically verified in **Figure 4A** and visualized in **Figure 5B**, the most striking result observed in **Figure 5C** is that, on average, vector strength is consistently greatest during spontaneous activity and consistently weakest during odor-driven responses; this pattern of results is consistent across all recordings and across both filtering ranges, with only one exception; in this case, vector strength values were roughly equal (see **Figure 5D**$_5$). Also consistent with **Figures 5A,B**, vector strength values during the late response epoch were, on average, lower than the spontaneous epoch but greater than the early response epoch.

Given that CSD analysis indicated that distributed oscillatory activity was in a higher frequency range during odor responses (**Figure 5D**), we might have expected to observe greater vector strength values during odor-driven responses in the higher frequency domain. However, as shown in **Figures 5A,D**, the relative disparity between spontaneous and odor-driven phase locking was not statistically different as a function of the filter implementation and we found no examples where odor-driven phase locking was, on average, stronger than spontaneous phase locking. Thus, whether considering phase locking at the unitary or population-level, or at low versus high frequency ranges, the results shown here are in contrast to what the transient oscillatory model predicts; namely that phase locking should be (relatively speaking) an odor-driven phenomenon.

Finally, comparison of mean unitary vector strength values (**Figure 5A**) with mean population-level vector strength values (**Figures 5C,D**) provides an indicator of how well synchronized the population is on an oscillatory time scale. Specifically the mean unitary vector strength should be generally higher than those calculated based on a population vector simply because variability in preferred phase angle is ignored when averaging unitary vector strength. Recall that vector strength is defined as the reciprocal of the circular variance of the distribution of phases of sampled spikes to the LFP; the preferred phase angle is simply the mean of a given distribution of spikes. Thus as the variability in spike phases increases, vector strength decreases. Given a set of individual units with different preferred phase angles, when treated as a population, those different preferred phase angles (i.e., different means) will add variance to the population-level distribution of spike phases, resulting in lower population-level vector strength. The ratio of the mean population vector strength to mean unitary vector strength indicates the relative ability of all of the neurons to spike in the same phase (*sic.* synchronously). Based on the 25 to 55-Hz LFPs, the population: unit vector strength ratio of the three response epochs was 0.63, 0.33, 0.48 for spontaneous, early, and late response epochs respectively. These results indicate that the preferred phase angles during spontaneous activity are more consistent (i.e., producing more synchronous spiking) than the response epochs within this LFP frequency range. For the 55 to 85-Hz LFPs, the ratios were 0.40, 0.61, and 0.61 for spontaneous, early, and late response epochs respectively suggesting that units were relatively less phase aligned spontaneously in this frequency range. This pattern of results suggests that synchronous spiking from the population is more likely in the higher frequency range. In all cases, however, there is variation in individual units' preferred phase relationships to the oscillations, which results in lower network synchronization. Furthermore, even within the higher LFPO frequency domain, vector strength values during the odor-driven response are still quite small.

THE PHASE RELATIONSHIP BETWEEN SPIKING AND OSCILLATIONS IS SMALL BUT SIGNIFICANT

Relative to vector strength values commonly described in various regions of the vertebrates brain such as the auditory system (e.g., Goldberg and Brown, 1969; Moushegian et al., 1975; Kadner and Berrebi, 2008), the average vector strength values observed in the present study as well as those from other insect studies (Ito et al., 2009) are relatively small. However, we can determine whether these values are above what is expected by chance, by

randomly jittering the spike timing within a restricted time range and recalculating vector strength. If the phase relation between spikes and the LFP are above what is expected by chance, then randomly jittering spike times will negatively impact vector strength values while maintaining much of the internal statistics of the dataset. First, to visualize this point, we randomly jittered all of the spikes from an individual spike train from a single recording ($N = 44{,}758$ across a 166-min recording session) by 0, ±3, or ±6 ms and cross correlated the spike train to one of the corresponding LFP recording sites (**Figure 6A**). Data includes both spontaneous and odor-driven activity. Note that in contrast to the cross correlations between LFPs (see **Figure 4**), this unit–LFP cross correlation shows clear temporal structure. Furthermore, random jittering of spikes by ±3 ms reduced and by ±6 ms eliminated this cross correlation. Next, using a bootstrapping method where spike trains were randomly re-jittered and vector strength recalculated, vector strength decay can be quantified as a function of the time range of the random jittering. In this case we used all units from a single recording and only spikes occurring within our three 300 ms peri-stimulus epochs ($N = 19$ units, 10,762 spikes). The initial data were re-jittered 30 times for each jitter width and vector strength was recalculated. Results of this analysis were then statistically analyzed using one-tailed paired *t*-tests. **Figure 6B** plots the mean vector strength by jitter width for each of the three peri-stimulus epochs and clearly indicate that the distributions of spikes tend to occur at a particular phase of the oscillation cycle. One-way *t*-test comparison of mean vector strength (significant results inset as asterisks) indicates a significant drop in strength as a function of increasing the jitter width by as little as ±1 ms up to the point at which the vector strength values approach 0 (±5 ms; $p < 0.001$; **Figure 6B**); this was true independent of the stimulus epoch used in the analysis. These results establish that while phase locking is weak relative to other sensory systems, it is nevertheless greater than what is expected based on chance.

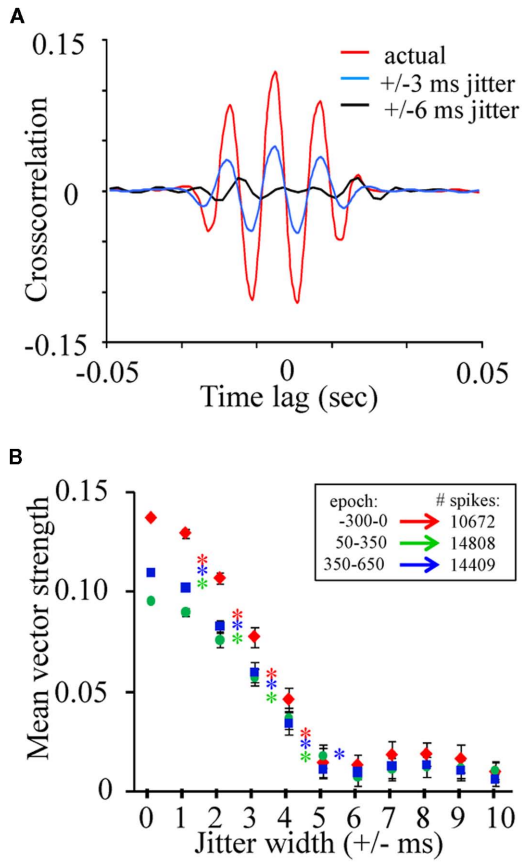

FIGURE 6 | The phase relationship between spiking and oscillations is significant. (A) The cross-correlation of a single unit to a parallel recorded LFPO normally (in red), after spikes were randomly jittered by a width of ±3ms (blue) and after spikes were jittered by a width of ±6 ms (black). **(B)** Vector strength for the three epochs and their recalculation after jittering. Jittering width started at ±1 ms and was increased in 1 ms steps. Inset asterisks indicate significant differences in vector strength to neighboring jittering ranges (one-way paired *t*-tests; $p < 0.01$). Note that the resulting strength of the unit–LFP relationship is clearly dependent by the degree of jittering indicating that these correlations, though small, are dependent on precise spike timing with respect to the periodic nature of the local field.

BICUCULLINE SIGNIFICANTLY REDUCES ODOR-DRIVEN FM OSCILLATORY RESPONSES AND DECREASES THE LATENCY OF SPIKING RESPONSES

Previous intracellular investigations on the effects of $GABA_A$ receptor blockade establish that BMI injection into *Manduca* ALs specifically and reversibly blocks a $GABA_A$-dependent inward Cl^- conductance in PNs (Christensen et al., 1998). This conductance normally results in a brief IPSP (44 ± 31 ms; Staudacher et al., 2009) and suppression of spontaneous PN spiking followed by a relatively prolonged supra-threshold EPSP, upon which an excitatory burst of spikes is superimposed. Thus, we first sought to establish if an additional function of $GABA_A$ receptors in the AL is to mediate the network's ability to maintain distributed oscillatory activity across the AL using the same methods as above. This analysis was performed using a within-animal design so that the results could be directly compared with those obtained prior to $GABA_A$ receptor disruption.

First, application of BMI disrupted and in many cases completely eliminated the odor-driven FM LFPOs. **Figure 7A** presents four typical TFRs from four different moths before and during BMI application. In all four cases, application of BMI caused a qualitative decrease in oscillatory power relative to pre-BMI measures. This decrease ranged from near complete (**Figure 7A** panels 1 and 2) to partial (panels 3 and 4). Note too that in **Figure 7A2**$_{\text{i}}$ the spontaneous 20–50 Hz activity, which terminates with odor stimulation, was also greatly reduced as compared to **Figure 7A2**$_{\text{ii}}$. This suggests that both ongoing and odor-driven oscillations were decreased. In order to statistically verify the loss of the FM responses, we modeled *z*-score normalized oscillatory power as a function of the main effects of individual differences between moths, the recording site the frequency step and the pharmacological treatment and their interactions using ANOVA. Results were based on five recordings for which we had comprehensive datasets for both pre and during BMI treatment. Results of the ANOVA indicate a significant drop in power from 0.32 ± 0.98 to −0.37 ± 0.81 (mean ± SD) as a result of BMI application ($p < 0.001$). Importantly the interaction of frequency by treatment was also significant. *Post hoc* analysis of oscillatory

power as a function of the frequency by treatment interaction was performed using one-tailed paired *t*-tests comparing power at each of 84 frequency steps a *p*-value of 0.0005 was used to reduce experiment-wise type 1 error rate (0.05/84). Results of this analysis (inset in **Figure 7B** indicated that oscillatory power at nearly every frequency step from 25 to 100 Hz was significantly reduced during BMI application.

The corresponding spiking responses were also impacted by BMI application. **Figure 7C** represents the mean peri-stimulus response to all presentations of all odors across all five moths used in this analysis. Note that as the oscillatory activity is decreased during BMI application (**Figures 7A,B**), there is a corresponding ∼60 ms decrease in the mean onset latency of the spiking population response (inset black arrow). There is also an increase in spike rate (inset gray arrow) and overall duration of the spiking response. This pattern of results is consistent with the previous findings of Christensen et al. (1998). That is, by disrupting the inward Cl^- conductance specifically, I_1 inhibition is presumably lost and hence, PNs respond earlier, stronger, and longer.

$GABA_A$ RECEPTOR BLOCKADE INCREASES ODOR-DRIVEN DISTRIBUTED ACTIVITY BUT CHANGES FREQUENCY CONTENT

A paired one-tailed *t*-test revealed that across all five moths and the six possible LFP–LFP cross correlations within each, $GABA_A$ receptor blockade did not significantly affect the mean peak LFP–LFP cross-correlation during spontaneous activity ($p = 0.28$) but did significantly increased the peak correlation for odor-driven epochs ($p < 0.01$; **Figure 8A**). Persistent cross correlations across recording sites during $GABA_A$ receptor blockade suggests that distributed activity is not entirely $GABA_A$-dependent. Comparison of the pre versus during BMI cross correlations shown in **Figure 8A**$_{ii}$ suggests that odor-driven distributed frequency content has been affected.

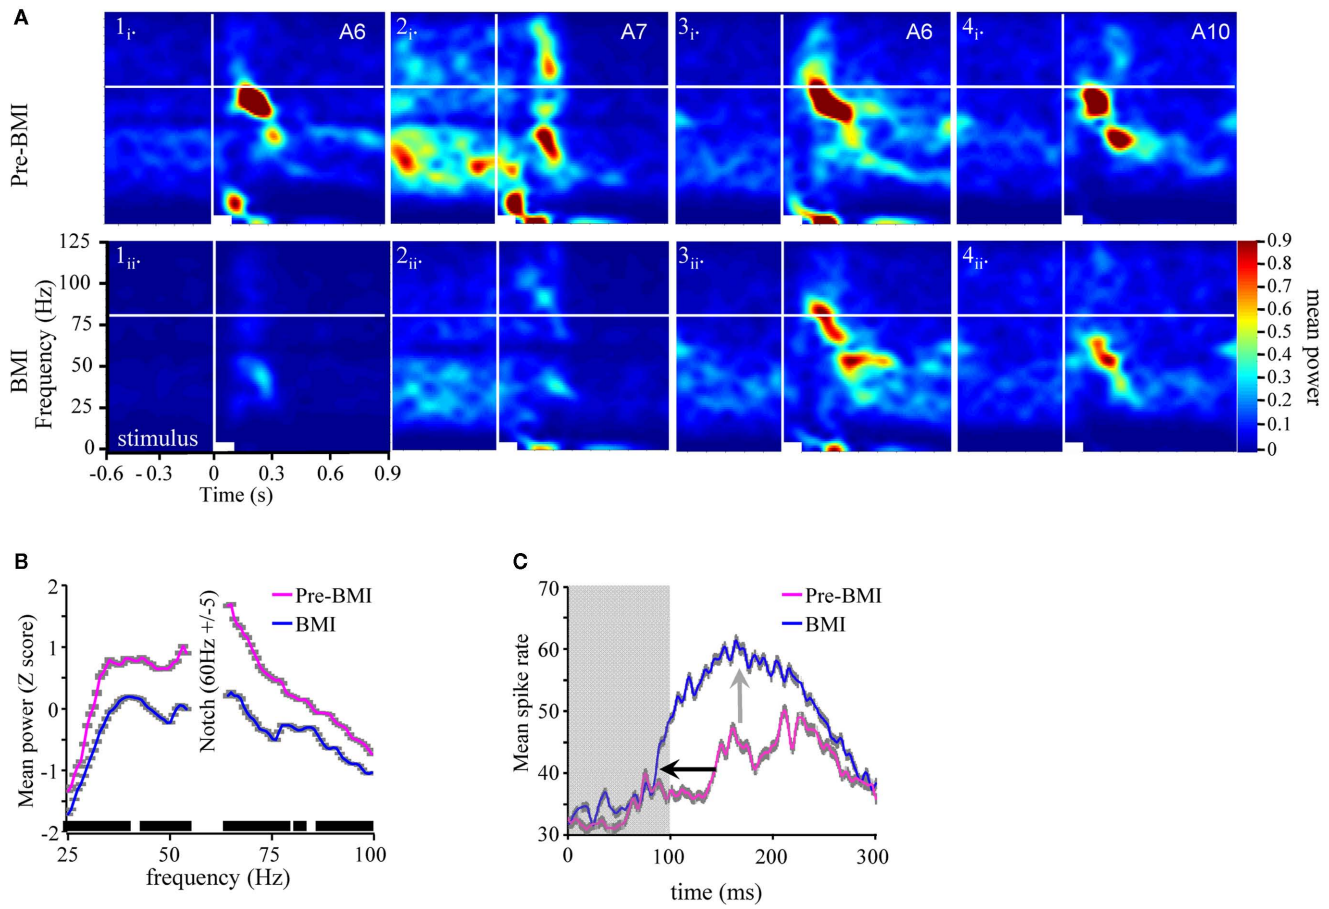

FIGURE 7 | Bicuculline significantly reduces odor-driven FM oscillatory responses (A). TFR analyses from four separate experiments (columns 1–4) in response to odor prior to **(i)** versus during BMI application **(ii)**. All panels are the result of TFR analyses that were calculated, normalized, and presented as described in **Figure 1**. Inset in the upper right corner of each panel **(i)** are the odorants used to generate both **(i)** and **(ii)**. **(B)** Mean *z*-score normalized oscillatory power by frequency for pre versus during bicuculline (BMI) application. Inset in gray shading are the standard errors for each frequency step. We also performed a *post hoc* one-tailed paired *t*-test. Significance level was set at 0.0005 to maintain an overall *post hoc* type 1 error rate of 0.05. Inset black bars above *X*-axis represent ranges of frequencies where power was significantly decreased as a function of BMI application; breaks in the bar indicate frequencies that were not statistically different. Overall only 16 of 84 tests were not significant. **(C)** Mean peri-stimulus population response histogram averaged over all responses to all odors for the five moths used in this series of analyses (gray shaded area around each trace represents SEM). Gray box indicates Stimulus duration. Inset arrows highlight the decrease in response onset latency (black) and increased spike rate during the response (gray).

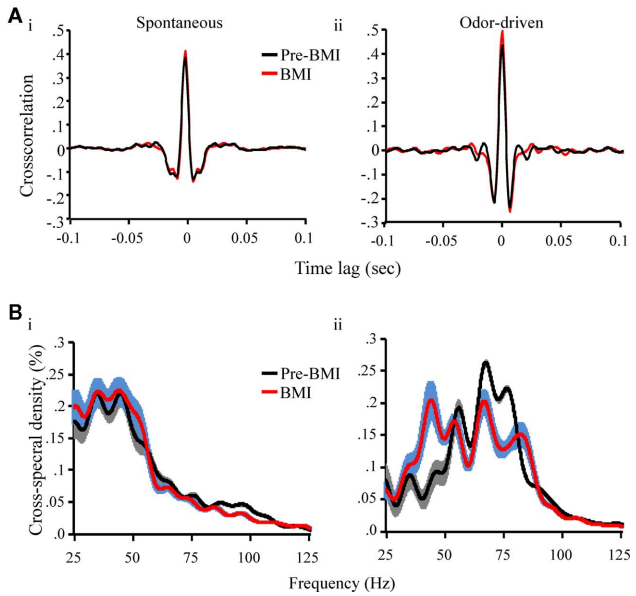

FIGURE 8 | GABA$_A$ receptor blockade increases odor-driven distributed activity and broadens its periodic structure (A). The mean cross correlation prior to BMI application (red) and during BMI application. Results displayed for spontaneous **(i)** and odor-driven **(ii)** epochs. Results based on all possible pair wise comparisons across the four LFP recording sites. **(B)** Mean cross spectral density prior to BMI application (red) and during BMI application. Results displayed for spontaneous **(i)** and odor-driven **(ii)** epochs and based on the data represented in **(A)**.

To address this question statistical comparison of CSD measures was performed using ANOVA. **Figure 8B** displays the mean CSD for spontaneous and odor-driven activity for pre versus during BMI application. ANOVA revealed that there was no significant change in mean CSD during spontaneous activity ($p > 0.01$) indicating that the relative power of distributed oscillations remained largely unchanged during this epoch. However, there was a significant shift in distributed frequency content during odor-driven responses that resulted in an increase in lower frequency content (between 40 and 55 Hz) while at the same time causing a relative reduction of higher frequency content (65–85 Hz; $p < 0.01$; **Figure 8B**).

GABA$_A$ RECEPTOR BLOCKADE DECREASES PHASE RELATIONSHIPS BETWEEN UNITARY SPIKING AND LOCAL FIELD POTENTIAL OSCILLATIONS

Finally, we investigated whether BMI application affected the phase relationship between unitary spiking and oscillatory activity. As mentioned above, results of the GLM analysis of vector strength indicated the BMI effect was significant ($p < 0.01$). Comparison of mean vector strength indicated that BMI application reduced mean vector strength (collapsed across all three peri-stimulus epochs) by 11% (pre-BMI = 0.14; BMI = 0.12). We again found that phase locking during spontaneous epochs was significantly stronger than either early or late odor-driven response epochs (**Figure 9A**). On a unit by unit basis BMI produced a notable decrease in phase locking in both spontaneous and odor-driven responses (**Figure 9B**). For purpose of comparison, **Figure 9B** data are based on the same units and in response to the same odor (1-decanol) as shown in **Figure 5B**. There was qualitative evidence of a phase relationship remaining in some cells in **Figure 9B**. However, as compared to pre-BMI application (see **Figure 5B**) phase locking has been greatly reduced during both spontaneous and odor-driven activity. Nevertheless, as demonstrated in **Figure 9A**, phase locking clearly remains across all recordings.

Figures 9C,D shows population-level vector strength as a function of peri-stimulus epoch for both the low (**Figure 9C**) and high (**Figure 9D**) band pass filter settings; these data too are from the same animals and hence directly comparable to **Figures 5C,D**. Here again, phase locking during spontaneous activity was consistently greater than either of the response epochs with only one exception (see **Figure 9D$_4$**). These results suggest that while phase locking is significantly reduced, what remains is still greatest during spontaneous rather than odor-driven activity.

DISCUSSION

Currently there is a long standing debate about the role of temporal processing in odor identity encoding, at the level of primary olfactory networks. One hypothesis, the transient oscillatory synchrony model, posits that odor stimulation drives LFPOs, which in turn drives precise spike time synchronization of projection neurons from across the AL (i.e., in a distributed manner), thereby binding this distributed output at the level of MB input. These synchronized spikes are the proposed information carriers for subsequent odor identification by the MB. The MB in turn is structurally and functionally organized to receive and interpret synchronous input from the AL on an oscillatory timescale (for review see Laurent, 2002). This is an exciting model because it provides several predictions, some of which we have attempted to test.

The first major finding of the present study is that odor stimulation drives oscillatory responses that are strongly frequency modulated. Furthermore, this modulation occurs in a stimulus-dependent manner. Among vertebrates, both spontaneous and stimulus driven oscillations have been observed across several brain regions (including the OB) and across a broad spectral range of frequencies (for review see Buzsaki and Draguhn, 2004). In the olfactory system of mammals, stimulus driven oscillations can be observed *in vivo* from 15 to 35 Hz (beta band) to 40–90 Hz (gamma band; Buonviso et al., 2003; Martin et al., 2004) and there is evidence that spiking can phase lock to these oscillations during a stimulus driven response (Kashiwadani et al., 1999). Odor-driven oscillations have typically been described as a constant narrow band (*sic* 20–30 Hz) in insects, such as the locust and honey bee (e.g., Laurent and Davidowitz, 1994; Wehr and Laurent, 1996; Stopfer et al., 1997). Hence, according to the transient oscillatory model these oscillations should be of a reasonably constant frequency that should be independent of the odorant presented (Laurent et al., 1996a). However, we find that in the AL of *Manduca*, as in vertebrates, oscillations occur across a far broader spectral range, modulating across response time from as high as ∼100 Hz down to as low as ∼20–30 Hz; this modulation takes no more than ∼400 ms to evolve given a 100-ms stimulus and does so in an odor-specific manner. Odor-driven FM oscillations and their tendency to occur in two distinct epochs within the AL, while until now

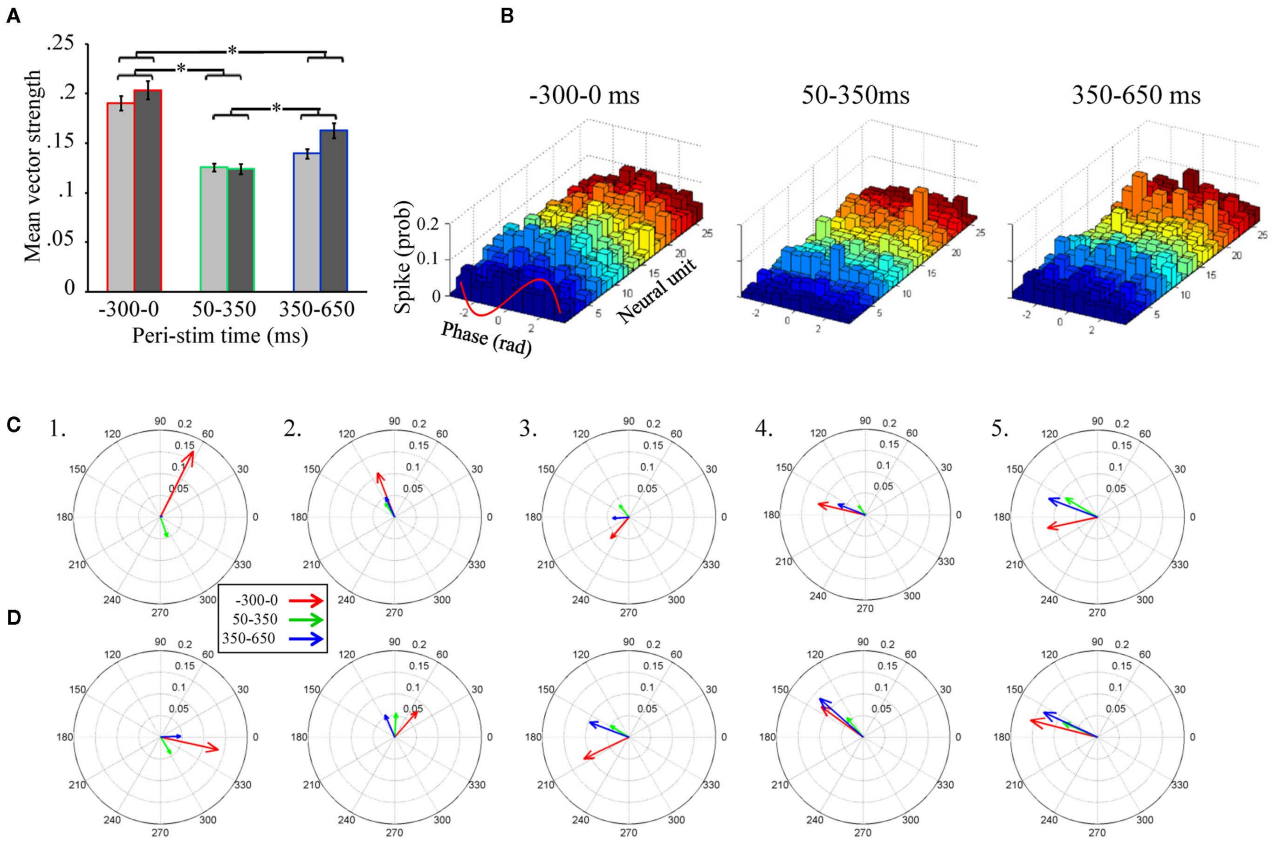

FIGURE 9 | GABA$_A$ receptor blockade decreases phase correlations between unitary spiking and local field oscillations. **(A)** Mean vector strength as a function of the three stimulus epochs during BMI application. Results are broken down by filter implementation (25–55 Hz in light gray; 55–85 Hz in dark gray). The results of statistical *post hoc* comparison between means are inset; means with significant differences indicate by an asterisk. Errors bars represent standard errors. **(B)** Phase lag histograms for 26 individual units from a single recording during spontaneous and the two odor-driven response epochs during BMI application. For comparative purposes, these histograms were generated from the same moth and in response to the same odorants as shown in **Figure 5B**. **(C)** Population-level measures of vector strength as a function of the three stimulus epochs for the five moths used in this analysis. Results are for the 25 to 55-Hz band pass filter implementation and each panel represents a different animal. **(D)** The same as **(C)**. except using the 55 to 85-Hz band pass filter.

unreported in the insect literature, is strikingly similar to what has been observed in mammalian OB (Buonviso et al., 2003).

It has recently been shown that prolonged odor stimulation (ca. 500–4000 ms) drives FM oscillations ranging from 50 Hz down to 10 Hz in the MBs of *Manduca* (Ito et al., 2009). We find that FM LFPOs recorded in the AL occur across a broader spectral range, do not require prolonged stimulation and are furthermore clearly odor dependent. The two studies taken together suggest that the highest frequency content does not pass from AL to MB. Indeed, much of the higher frequency content that occurs during the early phase of the AL response (100–80 Hz at ∼60–110 ms post stimulus onset) occurs prior to initiation of PN spiking responses, which typically start no earlier than ∼110–120 ms after the odor valve is actuated (Daly et al., 2004b; Staudacher et al., 2009). Thus, at least some of the higher frequency content observed in the AL, because it occurs while PN output is suppressed, cannot be transmitted to the MB.

In insects, it has been proposed that the downstream receivers of AL output, the Kenyon cells of the MB, are normally under strong inhibitory control by inputs projecting from the lateral horn. Lateral horn cells, like the Kenyon cells of the MB calyx, receive direct excitatory input from the AL. It has been furthermore proposed that this circuitous pathway establishes an integration window that is opened for synchronous excitatory input from the AL then abruptly closed by inhibitory input from the lateral horn (Perez-Orive et al., 2002). This integrate-and-reset window occurs within the periodic timescale initially established by odor-driven AL oscillations. From a theoretical perspective, given that neurons have fixed conduction velocities, cable lengths, and synaptic delays, the circuit function proposed by Perez-Orive et al. (2002) will have a biophysically constrained integration window. As a result, such an integrate-and-reset circuit will have limited tolerance for the FM oscillations that we have observed (for review see Laurent, 2002). It remains to be determined if and how subsequent decoding in the MB could occur under the constraints of a static integration window.

A second finding of this study is that odor-driven FM oscillatory responses are largely restricted to a subset (less than 2 on

average) of the four LFP recording sites and hence appears in most cases to be regionalized. Furthermore, cross correlation analysis also indicates that there is only a very small amount of activity that is both distributed across pairs of recording sites and temporally structured (i.e., oscillatory). This distributed oscillatory activity is state dependent; that is, spontaneous distributed activity is dominated by oscillatory activity in the 25 to 55-Hz range, whereas distributed odor-driven oscillations are in the 55 to 85-Hz range; this pattern of results of relatively higher frequency oscillatory content during odor-driven responses is consistent with prior findings in honeybee (Ritz et al., 2001). While the precise degree of this regionalization of responses remains an important area for further research, it is at odds with the transient oscillatory model.

One question that arises from these results is whether the relatively small amount of distributed oscillatory activity is sufficient to correlate with or otherwise influence spiking behavior in PNs in a manner that increases the amount of synchrony across AL PNs, as required by the transient oscillatory model. Our results demonstrate that spike timing is indeed biased to occur in a preferred phase relationship to the oscillation, though this bias was clearly state and cell (and recording) specific (see **Figures 5** and **6**). Contrary to the expectations of the transient oscillatory model, however, we observed that during an odor-driven response, as a population, this phase locking was largely lost. This occurred in part because unitary vector strength values dropped significantly during the odor response, which means that even at the level of individual cells, odor stimulation decoupled cells from the influence of LFPOs. Additionally, preferred phase angles of individual units became more variable and hence more evenly distributed across all phases of the oscillation cycle; this lowered population-level measures of phase locking. Thus, odor-driven oscillations *per se* had little effect on synchronizing the spikes across individual units of the recorded population. It is worth highlighting that strong phase locking has been typically observed in response to longer stimulations than we presented here (Laurent and Naraghi, 1994; Laurent et al., 1996b; Stopfer et al., 1997; Kashiwadani et al., 1999). While future work will explore the possibility that longer and temporally structured stimuli have differing effects on odor-driven phase locking, we note that 100 ms stimulations are entirely consistent with expected encounter times in natural plumes that the moth might encounter (Murlis and Jones, 1981) and are known to respond to in behavioral studies (Willis and Baker, 1984).

The fact that phase locking is lost during odor-driven responses in this model system should not be surprising. Indeed, it has long been known that in response to odor stimulation, AL PNs in *Manduca* are briefly inhibited followed by a sustained supra-threshold depolarization and spiking (for several examples see Christensen et al., 1996; Christensen et al., 1998; Heinbockel et al., 1998, 1999; Staudacher et al., 2009). This explains in part why odor responses appear to be characterized by a loss of phase relationship between our recorded units and LFPOs. This also likely explains why others were not able to establish phase locking during odor-driven responses in previous studies (Christensen et al., 2003). However, this does not explain the prolonged loss of phase locking.

In the AL of *Manduca*, LNs and PNs are spontaneously active under normal conditions; this has long been known (Matsumoto and Hildebrand, 1981; Kanzaki et al., 1989). Because oscillations are presumably the result of reciprocal synaptic connectivity within the network, it is expected that there is always some spontaneous oscillatory activity observable in the LFP, even though the amplitude may be small, relative to what is observed during odor-driven responses. Unexpectedly, we found that there is a tendency for individual cells to spontaneously spike in phase with oscillation cycles in the absence of odor stimulation.

It has been shown that even weak oscillatory signals in neural circuits can in some cases enhance signal to noise ratios; this phenomena is called stochastic resonance (Wiesenfeld and Moss, 1995; Wilkens et al., 2002; Korn and Faure, 2003). It has also been shown that weakly correlated noise in the olfactory circuit can enhance spike time synchrony (Galán et al., 2006). Perhaps then, the most intriguing finding in the current study is the significantly greater spike to LFP phase locking observed during spontaneous activity, relative to odor-driven responses; this was true independent of the frequency range of the oscillations under consideration. Thus, if oscillations were acting as a synchronizing mechanism in this model system, they are doing so during ongoing, spontaneous activity, and not during odor-driven responses. This is in sharp contrast to predictions made by the transient oscillatory model, which assumes that phase locking is a stimulus driven phenomenon involved in odor identity encoding. It is worth noting that we are unaware of any reported comparisons of spontaneous versus odor-driven phase locking prior to the present study, thus it is difficult to assess the generality of our findings. However, "bouts" of spontaneous oscillatory activity have been observed in both locust (Laurent and Naraghi, 1994) and honeybee (Ritz et al., 2001) olfactory pathways that appear to be consistent with our observations.

In response to odor stimulation, individual PNs recorded from *Manduca* AL are briefly inhibited via an inward Cl^- conductance, then burst as the second epoch of oscillations emerge. Several lines of evidence from both vertebrate and invertebrate models establish that spatially co-localized principal output neurons tend to produce stimulus driven synchronous spike bursting patterns (Schoppa and Westbrook, 2001, 2002; Lei et al., 2002; Daly et al., 2004b; Hayar et al., 2005) in what has been described as onset synchrony (Christensen et al., 2001; Lei et al., 2002). Parallel recordings of neural ensembles as well as serial reconstructions of AL output activity suggest that the onset of this burst of activity is different for different glomeruli thereby producing a brief sequence of onsets that are odor dependent and optimize within ∼240 ms of odor onset (or ∼120 ms from the onset of the excitatory response) in this model system (Daly et al., 2004b; Staudacher et al., 2009); this is roughly consistent with calcium-imaging studies in honeybee (Galán et al., 2004). This time course for producing odor-specific activation patterns also appears to be consistent with physiological evidence from other model organisms (Muller et al., 2002; Lehmkuhle et al., 2006; Spors et al., 2006; Namiki and Kanzaki, 2008; Namiki et al., 2009) as well as behavioral evidence demonstrating odor identification/discrimination within the same approximate time frame or faster (Laska et al., 1999; Uchida and Mainen, 2003; Budick and Dickinson, 2006; Wesson et al., 2008a,b). Thus, it is reasonable to conclude that the temporal window for odor identification

is in the range of 240 ms from stimulus onset or about ∼120 ms from initiation of the excitatory response within the AL (e.g., see **Figure 1**).

Stopfer et al. (1997) demonstrated that oscillations are abolished in the honeybee by application of the $GABA_A$ receptor antagonist picrotoxin. Based on a stimulus generalization paradigm they also suggest that $GABA_A$ receptor blockade disrupts fine odor discrimination in this insect. These findings represent a cornerstone of the transient oscillatory model because they identify a specific functional role for oscillations. More recently, however, the conclusion that $GABA_A$ receptor blockade affects discrimination only of closely related odors has been challenged by a more comprehensive behavior-pharmacological study in *Manduca* (Mwilaria et al., 2008). This study was based on pharmacological data from 1680 moths and established that discrimination of pairs of odors is generally disrupted independent of the relatedness of the odor pairs. Furthermore, Mwilaria et al. (2008) showed that $GABA_A$ receptor disruption increased detection thresholds. Based on these finding it was proposed that the disruption of the ability to discriminate closely related odors was most likely due to a general increase in detection thresholds, which increased task difficulty. These findings have been subsequently supported by physiological and wind tunnel experiments demonstrating that $GABA_A$ receptor blockade in the macroglomerular complex of the AL (the region of the male AL where the primary components of female pheromones are processed), does not affect pheromone detection at the level of spiking PNs but does reduce PN signal to noise (Lei et al., 2009). These researchers furthermore showed that reduction in signal to noise disrupts the ability of male moths to successfully plume track. This is consistent with the interpretation that at the level of sensory perception, loss of AL $GABA_A$ receptor function results in a loss of ability to clearly perceive the presence of (i.e., detect) an odor cue (Mwilaria et al., 2008). This too is in stark contrast to the transient oscillatory model, which predicts that the sole functional effect of $GABA_A$ receptor blockade is on the ability to discriminate closely related odors.

In conclusion, it is perhaps not surprising that the transient oscillatory model, which is based primarily on data from the locust, does not fit well with data from model systems such as *Manduca* given the differences in structural organization of the AL between these two species. For example, locusts have a derived AL morphology consisting of hundreds of small microglomeruli, in which both olfactory receptor neurons and PNs are multi glomerular and LNs are non-spiking (Anton and Homberg, 1999). In contrast, *Manduca,* like most insects and mammals, have exclusively uniglomerular ORNs and predominantly uniglomerular PNs as well as several LN morphologies, all of which produce action potentials. The differences in structure between the locust and *Manduca* likely reflect functional adaptations to distinct behavioral and chemical ecologies. Insects such as *Manduca* are more or less specialized upon a single food source as adults (i.e., floral nectars), and heavily dependent upon olfactory cues for locating these food sources from a distance (Bernays, 2001). For these insects there is a need to detect and rapidly discriminate among many faint and brief odor cues in the environment in order to locate a particular food source via odor plume tracking behavior. In contrast, orthopteran insects including locusts are extreme dietary generalists, consuming a wide variety of food sources consisting of both living and dead plant and animal material (Gangwere, 1961). Individual food preferences appear to be based upon a mixture of variables including local environmental conditions, the need to avoid predators (Despland and Simpson, 2005), what types of food are available, food quality, and nutrient need (Raubenheimer and Simpson, 2003). Given this pattern of food acquisition, it is likely that these insects are less dependent on identifying palatable food sources using specific olfactory cues, but may instead navigate to broad categories of food odors such as green leaf volatiles (Ochieng and Hansson, 1999; Chen and Kang, 2000). Unfortunately, while recent studies now support the notion that the locust can learn and discriminate between a pair of broadly different odorants (Simoes et al., 2011) there are no studies to date that integrate observations of fine odor discrimination, food choice, and AL physiology in locusts so it is not possible at present to speculate how this insect actually detects and utilizes odors in a behavioral context. We can, however, posit that the locust, an insect with an unusual wide open feeding ecology, in addition to possessing a unique AL anatomy and physiology, is unlikely to process odorants in a manner typical of the majority of insects or of odorant-sensing animals in general.

ACKNOWLEDGMENTS

We thank Dr's S. M. Farris, G. Spirou, B. G. Schreurs, and A. Agmon for critical review of this research. This research was funded by NIH/NIDCD R01DC009417 and NIH/NCRR P20RR015574 to Kevin C. Daly and by the Mt. Sinai Health Care Foundation and the Alfred P. Sloan Foundation to Roberto F. Galán.

REFERENCES

Anton, S., and Homberg, U. (1999). *Insect Olfaction.* Berlin, NY: Springer, 457.

Bell, R. A., and Joachim, F. G. (1976). Techniques for rearing laboratory colonies of tobacco hornworms and pink bollworms. *Ann. Entomol. Soc. Am.* 69, 365–372.

Bernays, E. A. (2001). Neural limitations in phytophagous insects: implications for diet breadth and evolution of host affiliation. *Annu. Rev. Entomol.* 46, 703–727.

Boser, B. E., Guyon, I. M., and Vapnik, V. N. (1992). "A training algorithm for optimal margin classifiers," in *COLT 1992 Proceedings of the Fifth Annual Workshop on Computational Learning* (New York: ACM).

Budick, S. A., and Dickinson, M. H. (2006). Free-flight responses of *Drosophila melanogaster* to attractive odors. *J. Exp. Biol.* 209, 3001–3017.

Buonviso, N., Amat, C., Litaudon, P., Roux, S., Royet, J. P., Farget, V., and Sicard, G. (2003). Rhythm sequence through the olfactory bulb layers during the time window of a respiratory cycle. *Eur. J. Neurosci.* 17, 1811–1819.

Buzsaki, G., and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science* 304, 1926–1929.

Chen, H. H., and Kang, L. (2000). Olfactory responses of two species of grasshoppers to plant odours. *Entomol. Exp. Appl.* 95, 129–134.

Christensen, T. A., D'Alessandro, G., Lega, J., and Hildebrand, J. G. (2001). Morphometric modeling of olfactory circuits in the insect antennal lobe: I. Simulations of spiking local interneurons. *Biosystems* 61, 143–153.

Christensen, T. A., Heinbockel, T., and Hildebrand, J. G. (1996). Olfactory information processing in the brain: encoding chemical and temporal features of odors. *J. Neurobiol.* 30, 82–91.

Christensen, T. A., Lei, H., and Hildebrand, J. G. (2003). Coordination of central odor representations

through transient, non-oscillatory synchronization of glomerular output neurons. *Proc. Natl. Acad. Sci. U.S.A.* 100, 11076–11081.

Christensen, T. A., Waldrop, B. R., and Hildebrand, J. G. (1998). Multitasking in the olfactory system: context-dependent responses to odors reveal dual GABA-regulated coding mechanisms in single olfactory projection neurons. *J. Neurosci.* 18, 5999–6008.

Daly, K. C., Chandra, S., Durtschi, M. L., and Smith, B. H. (2001). The generalization of an olfactory-based conditioned response reveals unique but overlapping odour representations in the moth *Manduca sexta*. *J. Exp. Biol.* 204, 3085–3095.

Daly, K. C., Christensen, T. A., Lei, H., Smith, B. H., and Hildebrand, J. G. (2004a). Learning modulates the ensemble representations for odors in primary olfactory networks. *Proc. Natl. Acad. Sci. U.S.A.* 101, 10476–10481.

Daly, K. C., Wright, G. A., and Smith, B. H. (2004b). Molecular features of odorants systematically influence slow temporal responses across clusters of coordinated antennal lobe units in the moth *Manduca sexta*. *J. Neurophysiol.* 92, 236–254.

Despland, E., and Simpson, S. J. (2005). Food choices of solitarious and gregarious locusts reflect cryptic and aposematic antipredator strategies. *Anim. Behav.* 69, 471–479.

Friedrich, R. W., and Laurent, G. (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science* 291, 889–894.

Galán, R. F., Fourcaud-Trocme, N., Ermentrout, G. B., and Urban, N. N. (2006). Correlation-induced synchronization of oscillations in olfactory bulb neurons. *J. Neurosci.* 26, 3646–3655.

Galán, R. F., Sachse, S., Galizia, C. G., and Herz, A. V. M. (2004). Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput.* 16, 999–1012.

Gangwere, S. K. (1961). A monograph on food selection in Orthoptera. *Trans. Am. Entomol. Soc.* 87, 67–230.

Gelperin, A., and Tank, D. W. (1990). Odour-modulated collective network oscillations of olfactory interneurons in a terrestrial mollusc. *Nature* 345, 437–440.

Goldberg, J. M., and Brown, P. B. (1969). Response of binaural neurons of dog superior olivary complex to dichotic tonal stimuli: some physiological mechanisms of sound localization. *J. Neurophysiol.* 32, 613–636.

Hayar, A., Shipley, M. T., and Ennis, M. (2005). Olfactory bulb external tufted cells are synchronized by multiple intraglomerular mechanisms. *J. Neurosci.* 25, 8197–8208.

Heinbockel, T., Christensen, T. A., and Hildebrand, J. G. (1999). Temporal tuning of odor responses in pheromone-responsive projection neurons in the brain of the sphinx moth *Manduca sexta*. *J. Comp. Neurol.* 409, 1–12.

Heinbockel, T., Kloppenburg, P., and Hildebrand, J. G. (1998). Pheromone-evoked potentials and oscillations in the antennal lobes of the sphinx moth *Manduca sexta*. *J. Comp. Physiol. A* 182, 703–714.

Hosler, J. S., Buxton, K. L., and Smith, B. H. (2000). Impairment of olfactory discrimination by blockade of GABA and nitric oxide activity in the honey bee antennal lobes. *Behav. Neurosci.* 114, 514–525.

Ito, I., Bazhenov, M., Ong, R. C., Raman, B., and Stopfer, M. (2009). Frequency transitions in odor-evoked neural oscillations. *Neuron* 64, 692–706.

Johnson, B. A., Farahbod, H., Saber, S., and Leon, M. (2005). Effects of functional group position on spatial representations of aliphatic odorants in the rat olfactory bulb. *J. Comp. Neurol.* 483, 192–204.

Johnson, B. A., Woo, C. C., Hingco, E. E., Pham, K. L., and Leon, M. (1999). Multidimensional chemotopic responses to n-aliphatic acid odorants in the rat olfactory bulb. *J. Comp. Neurol.* 409, 529–548.

Johnson, B. A., Woo, C. C., and Leon, M. (1998). Spatial coding of odorant features in the glomerular layer of the rat olfactory bulb. *J. Comp. Neurol.* 393, 457–471.

Kadner, A., and Berrebi, A. S. (2008). Encoding of temporal features of auditory stimuli in the medial nucleus of the trapezoid body and superior paraolivary nucleus of the rat. *Neuroscience* 151, 868–887.

Kanzaki, R., Arbas, E. A., Strausfeld, N. J., and Hildebrand, J. G. (1989). Physiology and morphology of projection neurons in the antennal lobe of the male moth *Manduca sexta*. *J. Comp. Physiol. A* 165, 427–453.

Kashiwadani, H., Sasaki, Y. F., Uchida, N., and Mori, K. (1999). Synchronized oscillatory discharges of mitral/tufted cells with different molecular receptive ranges in the rabbit olfactory bulb. *J. Neurophysiol.* 82, 1786–1792.

Korn, H., and Faure, P. (2003). Is there chaos in the brain? II. Experimental evidence and related models. *C. R. Biol.* 326, 787–840.

Laska, M., Galizia, C. G., Giurfa, M., and Menzel, R. (1999). Olfactory discrimination ability and odor structure-activity relationships in honeybees. *Chem. Senses* 24, 429–438.

Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nat. Rev. Neurosci.* 3, 884–895.

Laurent, G., and Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science* 265, 1872–1875.

Laurent, G., and Naraghi, M. (1994). Odorant-induced oscillations in the mushroom bodies of the locust. *J. Neurosci.* 14, 2993–3004.

Laurent, G., Stopfer, M., Friedrich, R. W., Rabinovich, M. I., Volkovskii, A., and Abarbanel, H. D. I. (2001). Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu. Rev. Neurosci.* 24, 263–297.

Laurent, G., Wehr, M., and Davidowitz, H. (1996a). Temporal representations of odors in an olfactory network. *J. Neurosci.* 16, 3837–3847.

Laurent, G., Wehr, M., MacLeod, K., Stopfer, M., Leitch, B., and Davidowitz, H. (1996b). Dynamic encoding of odors with oscillating neuronal assemblies in the locust brain. *Biol. Bull.* 191, 70–75.

Lehmkuhle, M. J., Normann, R. A., and Maynard, E. M. (2006). Trial-by-trial discrimination of three enantiomer pairs by neural ensembles in mammalian olfactory bulb. *J. Neurophysiol.* 95, 1369–1379.

Lei, H., Christensen, T. A., and Hildebrand, J. G. (2002). Local inhibition modulates odor-evoked synchronization of glomerulus-specific output neurons. *Nat. Neurosci.* 5, 557–565.

Lei, H., Riffell, J. A., Gage, S. L., and Hildebrand, J. G. (2009). Contrast enhancement of stimulus intermittency in a primary olfactory network and its behavioral significance. *J. Biol.* 8, 21.

Linster, C., Johnson, B. A., Yue, E., Morse, A., Xu, Z., Hingco, E. E., Choi, Y., Choi, M., Messiha, A., and Leon, M. (2001). Perceptual correlates of neural representations evoked by odorant enantiomers. *J. Neurosci.* 21, 9837–9843.

MacLeod, K., and Laurent, G. (1996). Distinct mechanisms for synchronization and temporal patterning of odor-encoding neural assemblies. *Science* 274, 976–979.

Martin, C., Gervais, R., Chabaud, P., Messaoudi, B., and Ravel, N. (2004). Learning-induced modulation of oscillatory activities in the mammalian olfactory system: the role of the centrifugal fibres. *J. Physiol. Paris* 98, 467–478.

Matsumoto, S. G., and Hildebrand, J. G. (1981). Olfactory mechanisms in the moth *Manduca sexta*: response characteristics and morphology of central neurons in the antennal lobe. *Proc. R. Soc. Lond. B Biol. Sci.* 213, 249–277.

Moushegian, G., Rupert, A. L., and Gidda, J. S. (1975). Functional characteristics of superior olivary neurons to binaural stimuli. *J. Neurophysiol.* 38, 1037–1048.

Muller, D., Abel, R., Brandt, R., Zockler, M., and Menzel, R. (2002). Differential parallel processing of olfactory information in the honeybee, *Apis mellifera* L. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 188, 359–370.

Murlis, J., and Jones, C. D. (1981). Fine-scale structure of odor plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiol. Entomol.* 6, 71–86.

Mwilaria, E. K., Ghatak, C., and Daly, K. C. (2008). Disruption of GABA(A) in the insect antennal lobe generally increases odor detection and discrimination thresholds. *Chem. Senses* 33, 267–281.

Namiki, S., Haupt, S. S., Kazawa, T., Takashima, A., Ikeno, H., and Kanzaki, R. (2009). Reconstruction of virtual neural circuits in an insect brain. *Front. Neurosci.* 3:206–213. doi:10.3389/neuro.01.028.2009

Namiki, S., and Kanzaki, R. (2008). Reconstructing the population activity of olfactory output neurons that innervate identifiable processing units. *Front. Neural Circuits* 2:1. doi:10.3389/neuro.04.001.2008

Nusser, Z., Kay, L. M., Laurent, G., Homanics, G. E., and Mody, I. (2001). Disruption of GABA(A) receptors on GABAergic interneurons leads to increased oscillatory power in the olfactory bulb network. *J. Neurophysiol.* 86, 2823–2833.

Ochieng, S. A., and Hansson, B. S. (1999). Responses of olfactory receptor neurones to behaviourally important odours in gregarious and solitarious desert locust, *Schistocerca gregaria*. *Physiol. Entomol.* 24, 18–36.

Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science* 297, 359–365.

Raubenheimer, D., and Simpson, S. J. (2003). Nutrient balancing in grasshoppers: behavioural and physiological correlates of dietary breadth. *J. Exp. Biol.* 206, 1669–1681.

Ritz, R., Galán, R. F., Szyszka, P., and Herz, A. V. M. (2001). Analysis of odor processing in the mushroom bodies of the honeybee. *Neurocomputing* 38–40, 313–318.

Schoppa, N. E., and Westbrook, G. L. (2001). Glomerulus-specific synchronization of mitral cells in the olfactory bulb. *Neuron* 31, 639–651.

Schoppa, N. E., and Westbrook, G. L. (2002). AMPA autoreceptors drive correlated spiking in olfactory bulb glomeruli. *Nat. Neurosci.* 5, 1194–1202.

Simoes, P., Ott, S. R., and Niven, J. E. (2011). Associative olfactory learning in the desert locust, *Schistocerca gregaria. J. Exp. Biol.* 214, 2495–2503.

Spors, H., Wachowiak, M., Cohen, L. B., and Friedrich, R. W. (2006). Temporal dynamics and latency patterns of receptor neuron input to the olfactory bulb. *J. Neurosci.* 26, 1247–1259.

Staudacher, E. M., Huetteroth, W., Schachtner, J., and Daly, K. C. (2009). A 4-dimensional representation of antennal lobe output based on an ensemble of characterized projection neurons. *J. Neurosci. Methods* 180, 208–223.

Stevenson, R., Corbo, K., Baca, L., and Le, Q. (1995). Cage size and flight speed of the tobacco hawk moth *Manduca sexta. J. Exp. Biol.* 198, 1665–1672.

Stopfer, M., Bhagavan, S., Smith, B. H., and Laurent, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390, 70–74.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.

Uchida, N., and Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat. Neurosci.* 6, 1224–1229.

Waldrop, B., Christensen, T. A., and Hildebrand, J. G. (1987). GABA-mediated synaptic inhibition of projection neurons in the antennal lobes of the sphinx moth, *Manduca sexta. J. Comp. Physiol. A* 161, 23–32.

Wehr, M., and Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384, 162–166.

Wehr, M., and Laurent, G. (1999). Relationship between afferent and central temporal patterns in the locust olfactory system. *J. Neurosci.* 19, 381–390.

Wesson, D. W., Carey, R. M., Verhagen, J. V., and Wachowiak, M. (2008a). Rapid encoding and perception of novel odors in the rat. *PLoS Biol.* 6, e82. doi:10.1371/journal.pbio.0060082

Wesson, D. W., Donahou, T. N., Johnson, M. O., and Wachowiak, M. (2008b). Sniffing behavior of mice during performance in odor-guided tasks. *Chem. Senses* 33, 581–596.

Wiesenfeld, K., and Moss, F. (1995). Stochastic resonance and the benefits of noise – from ice ages to crayfish and squids. *Nature* 373, 33–36.

Wilkens, L. A., Hofmann, M. H., and Wojtenek, W. (2002). The electric sense of the paddlefish: a passive system for the detection and capture of zooplankton prey. *J. Physiol. Paris* 96, 363–377.

Willis, M. A., and Baker, T. C. (1984). Effects of intermittent and continuous pheromone stimulation on the flight behaviour of the oriental fruit moth *Grapholita molesta. Physiol. Entomol.* 9, 341–358.

Wright, G. A., Skinner, B. D., and Smith, B. H. (2002). Ability of honeybee, *Apis mellifera*, to detect and discriminate odors of varieties of canola (*Brassica rapa* and *Brassica napus*) and snapdragon flowers (*Antirrhinum majus*). *J. Chem. Ecol.* 28, 721–740.

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.

Received: 07 July 2011; paper pending published: 16 August 2011; accepted: 30 September 2011; published online: 25 October 2011.

Citation: Daly KC, Galán RF, Peters OJ and Staudacher EM (2011) Detailed characterization of local field potential oscillations and their relationship to spike timing in the antennal lobe of the moth Manduca sexta. Front. Neuroeng. ***4****:12. doi: 10.3389/fneng.2011.00012*

APPENDIX

METHODS

Discriminant analysis of time-frequency response spectrograms

The aim of this analysis was to establish whether odor-driven frequency modulation of the LFP oscillation was odor specific. There are two main approaches to investigate this: by means of unsupervised learning algorithms (e.g., k-means clustering analysis), or by means of supervised learning algorithms (e.g., discriminant analysis via support vector classifiers). In the unsupervised learning approach, all samples (in our case TFRs) are represented as points in a multidimensional space and groups of points that are closer among themselves than among others are classified together. One thus expects that all points from the same cluster belong to the same odor. In practice, however, this is seldom the case. The clouds of points from the same odor may not have radial symmetry and even worse, they may partially overlap with the "clouds" of other odors. Thus, the relative distance between points is not a good criterion to classify odors, regardless of the definition of distance (Euclidean, angular, etc.). The supervised classification approach, which is an alternative to unsupervised clustering, has been used previously to identify odor-specific patterns of neural activity in the antennal lobe of the honeybee (Galán et al., 2004). In this case, a multidimensional representation is also used but the points are labeled, i.e., assigned to an odor. Then, an optimization algorithm, the support vector classifier, (Boser et al., 1992) attempts to find a hyperplane (a plane in more than three dimensions) or another high-dimensional manifold that separates the points belonging to the same class (odor) from the rest. If successful, the manifold can be used as a classifier to identify that odor: on one side lie the points of the same odor; on the other side lie the rest. The same procedure can be repeated for each odor to compute odor-specific classifiers. Support vector classifiers have the advantage over clustering methods of tolerating some overlap (soft margin) between clouds of points from different stimuli. However, when the clouds overlap the mere existence of a separating manifold is not very informative, since both groups are not 100% separable. In such cases, the Classification Performance Index (CPI) provides a better parameter to quantify odor-specific TFR patterns; CPI is computed using the leave-one-out method (Boser et al., 1992; Galán et al., 2004): First, the separating manifold is computed after removing one point from the data set. Then, one tests if the point that was left out is correctly classified. These three steps (removal of a point, calculation of the discriminator, and classification of the point removed) are then iterated for all points of a given stimulus. The fraction of points that are correctly classified for each stimulus is the CPI. A high CPI means that the separating manifold is fairly insensitive to the removal of any given point, and hence robust to perturbations of the data set. A high CPI implies that the high-dimensional space in which the points are represented is divided into stimulus specific regions, despite some overlap between the clouds of points.

Dimensionality reduction

As a preprocessing step for TFRs discriminate analysis, we cropped out a reduced frequency–time window, which contained the vast bulk of the FM response across all TFRs (from 12 to 124 Hz and from −30 to 1000 ms); this resulted in a reduced data matrix of 118,800 pixels per TFR/response. For each response to each odor, we then defined a region of interest (ROI), consisting of all pixels that are above the 99th percentile of power density. All these ROI (9 odors × 20 repeats = 180 ROI) where overlaid to create a mask whose pixel values were "1" if that pixel was significantly activated by any odor in any trial and "0" if not. For the discriminant analysis with support vector classifiers, we only considered the pixels of the TFRs whose values in the mask were "1." Each TFR is thus represented as a vector 24,969 components (pixels) instead of 118,800 for the cropped TFR, thereby leading to a substantial dimensionality reduction of roughly 80%.

Kernels for the support vector classifier

The vectors representing the TFRs are fed into the algorithm of the support vector classifier. The output of the algorithms returns a set of n "support vectors," $\vec{s}_i$, weights a_i, and bias b that are used to classify a given vector $\vec{x}$ according to the following equation:

$$c = \sum_{i=1}^{n} a_i K(\vec{s}_i, \vec{x}) + b, \tag{A1}$$

where K is a *kernel* function. In the case of a linear kernel, it is the dot product: $K(\vec{s}_i, \vec{x}) = \vec{s}_i \cdot \vec{x}$ and the Eq. A1 defines a plane in the high-dimensional space. If $c \geq 0$, then $\vec{x}$ is classified as a member of group 1 (e.g., odor Y), otherwise it is classified as a member group 2 (e.g., any other odor different from Y). In this paper, we also use a non-linear kernel, specifically, a third order polynomial given by:

$$K(\vec{s}_i, \vec{x}) = (\vec{s}_i \cdot \vec{x} + 1)^3 .$$

This allows us to separate odors with a curvy manifold and the region assigned to a given odor may be composed of disjoint subregions. Using this kernel, we can discriminate TRFs from each odor group with 100% CPI.

Calculation of the confidence interval for the classification performance index

The null hypothesis to compute the confidence interval is that the two groups are completely intermingled and cannot be discriminated so that the probability that any point falls in one side of the classifier is $p = 50\%$. In our case, the two groups correspond to the $N = 20$ trials of the same odor (first group) and the rest of points from all other odors (second group). Thus the probability that M out of the N points fall into the same side of the classifier is given by a binomial distribution whose mean and variance are $N \times P$ and $N \times P \times (1\text{-}P)$, respectively. The binomial distribution itself is discrete. However, it can be smoothened by interpolation so that its cumulative distribution is also smooth, not staggered, and the percentiles can be calculated at any level. For example, the 95th percentile corresponds to the ratio $M/N = 66.25\%$. This means that if the CPI is above 66.25%, we can reject the null hypothesis with 95% confidence and for each CPI value above this level the p-value is smaller than 0.05. In the CPI plots, the red line indicates this significance level.

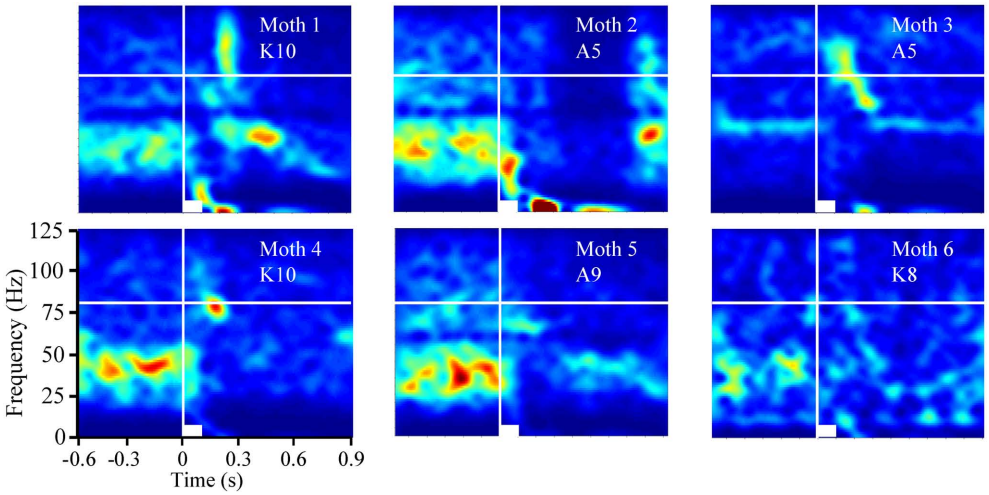

FIGURE A1 | Averaged time frequency response (TFR) spectrogram from 6 of 10 total recordings. Displayed in each panel are the averaged results of 20 independent analyses for each of the 20 presentations for various odors (inset; K = ketone, A = alcohol; number = carbon chain length). Also inset are white bars indicating stimulus duration (100 ms). Vertical white line indicates odor onset. Horizontal white line is an 80-Hz frequency reference. Note the ongoing oscillatory activity present centered around 30–50 Hz that occurs between 600 ms until just after stimulus onset. Upon stimulus onset this activity abruptly terminates. Note too that the weakest examples are moths three and six.

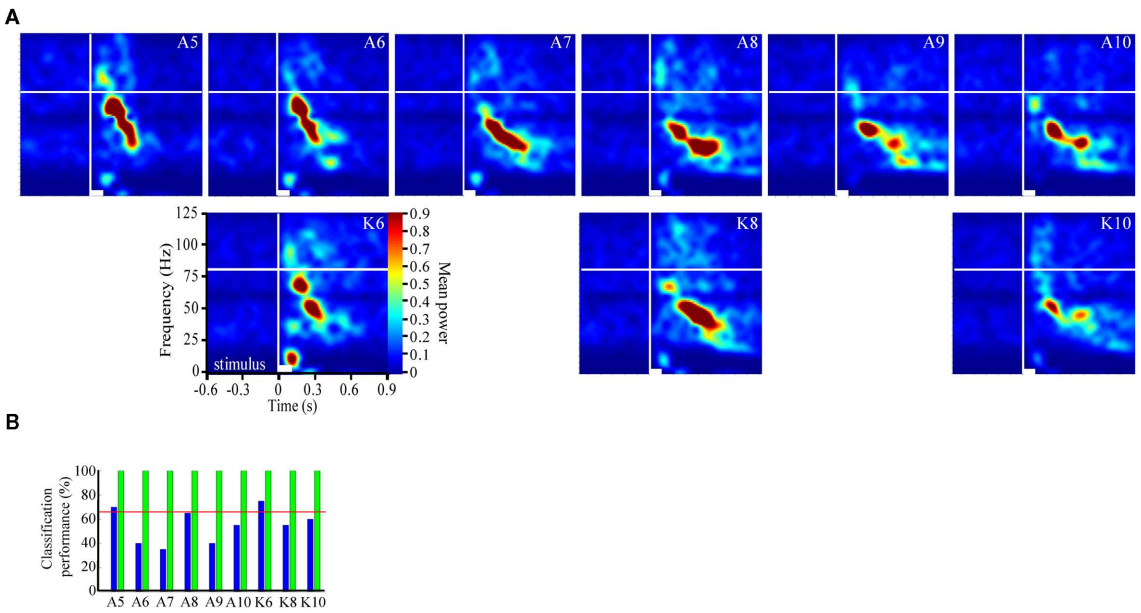

FIGURE A2 | Replication of Figure 2 using data from a different moth to highlight that FM modulation is odor dependent. (A) Each panel represents the averaged result of 20 individual TFR analyses that were calculated, normalized and presented as described in **Figure 1**. Note that alcohols and ketones of common length are aligned into columns. **(B)** Results of the discriminant analysis for the data shown in **(A)**. Blue Bars represent classification performance for the linear kernel function, whereas green bars represent performance of the third order polynomial kernel function. Inset red line represents the 95% confidence interval. Note that the linear kernel function classified responses of two of the nine odorants significantly above chance, whereas the polynomial function accurately classified all odor responses with 100% accuracy.

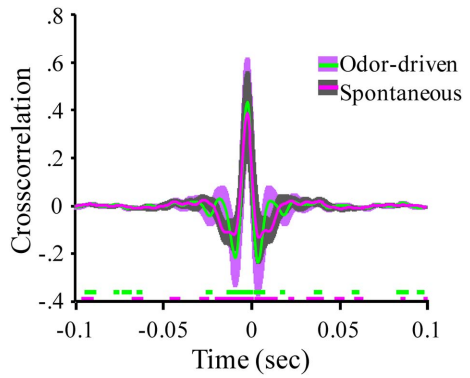

FIGURE A3 | Expanded view of the mean cross correlations for spontaneous versus odor-driven periods. Results based on all possible pair wise comparisons between the four LFP recording sites (six in total). Odor-driven cross-correlation based on 1 s samples, starting at odor onset, for each of 180 odor stimuli per moth ($N = 7$ moths). Spontaneous cross correlation based on 180 1 s samples taken immediately prior to odor onset for each odor stimulus. Note that cross correlations were calculated for each stimulus and results were averaged across seven moths. Shaded regions around the mean represent ± 1 SD. Inset broken lines just above the *X*-axis indicate where in time the cross correlation was significantly different from 0. Significance threshold for these tests was set to 2.3×10^{-6} to maintain an overall alpha of 0.01. Each color coded line corresponds to regions in time where the correlation was significantly different from zero for spontaneous (mauve) and odor-driven (green) cross correlations.

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 28 December 2011
doi: 10.3389/fneng.2011.00017

Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee

***Michael Schmuker*[1,2,3]*, *Nobuhiro Yamagata*[2,4†], *Martin Paul Nawrot*[1,3] *and Randolf Menzel*[2,3]**

[1] *Neuroinformatics and Theoretical Neuroscience, Institute of Biology, Freie Universität Berlin, Berlin, Germany*
[2] *Neurobiology, Institute of Biology, Freie Universität Berlin, Berlin, Germany*
[3] *Bernstein Center for Computational Neuroscience Berlin, Berlin, Germany*
[4] *Graduate School of Life Sciences, Tohoku University, Sendai, Japan*

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Roberto Fernández Galán, Case Western Reserve University, USA
Pawel Andrzej Herman, Royal Institute of Technology, Sweden

***Correspondence:**
Michael Schmuker, Neuroinformatics and Theoretical Neuroscience, Institute of Biology, Freie Universität Berlin, Königin-Luise-Str. 1-3, 14195 Berlin, Germany.
e-mail: m.schmuker@fu-berlin.de

†Present address:
Nobuhiro Yamagata, Behavioral Genetics, Max Planck Institute of Neurobiology, Martinsried, Germany.

The honeybee *Apis mellifera* has a remarkable ability to detect and locate food sources during foraging, and to associate odor cues with food rewards. In the honeybee's olfactory system, sensory input is first processed in the antennal lobe (AL) network. Uniglomerular projection neurons (PNs) convey the sensory code from the AL to higher brain regions via two parallel but anatomically distinct pathways, the lateral and the medial antenno-cerebral tract (l- and m-ACT). Neurons innervating either tract show characteristic differences in odor selectivity, concentration dependence, and representation of mixtures. It is still unknown how this differential stimulus representation is achieved within the AL network. In this contribution, we use a computational network model to demonstrate that the experimentally observed features of odor coding in PNs can be reproduced by varying lateral inhibition and gain control in an otherwise unchanged AL network. We show that odor coding in the l-ACT supports detection and accurate identification of weak odor traces at the expense of concentration sensitivity, while odor coding in the m-ACT provides the basis for the computation and following of concentration gradients but provides weaker discrimination power. Both coding strategies are mutually exclusive, which creates a tradeoff between detection accuracy and sensitivity. The development of two parallel systems may thus reflect an evolutionary solution to this problem that enables honeybees to achieve both tasks during bee foraging in their natural environment, and which could inspire the development of artificial chemosensory devices for odor-guided navigation in robots.

Keywords: dual pathway odor coding, mixture coding, antennal lobe, computational model, honeybee foraging

INTRODUCTION

Parallel olfactory subsystems are common in vertebrate and insect nervous systems. As Galizia and Rössler (2010) point out, two categories of parallel subsystems can be distinguished in the insect realm: segregate and dual parallel systems. Segregate parallel systems process different chemical stimuli (e.g., pheromone subsystems). In contrast, dual parallel systems analyze the same odorants, but with different coding and processing strategies. Often, such functional distinction is also evident in an anatomical separation.

A prominent case of a dual parallel olfactory system is found in the order of *Hymenoptera*, to which bees and ants belong. In these animals, uniglomerular projection neurons (PNs) relay olfactory information via two anatomically distinct tracts from the antennal lobe (AL) to higher brain centers, that is, the mushroom body (MB) and the lateral horn (LH), where sensory pathways from multiple modalities converge (see Galizia and Rössler, 2010 for a review). Those two tracts are called the lateral and the medial antenno-cerebral tract (l- and m-ACT), after their anatomical location in the brain (Mobbs, 1982). These pathways are specific to the group of hymenoptera, and hence they are not present in, for example, *Diptera* (flies) or *Orthoptera* (locust; Galizia and Rössler, 2010).

In the honeybee *Apis mellifera*, uniglomerular PNs sending their axons along either pathway innervate segregated populations of glomeruli, the basic functional units in the AL (Abel et al., 2001). The separation of those groups of PNs on the glomerular level may indicate that the information they convey is processed separately (Galizia and Rössler, 2010). The projections of PNs from both tracts target adjacent and partially overlapping regions in the MB and LH, suggesting that computational processes in those areas use information from both tracts (Müller et al., 2002; Kirschner et al., 2006).

Several studies have addressed functional differences in odor representation across the two tracts. Müller et al. (2002) reported that the segregation into two pathways is not related to the distinction between different odors or odors and pheromones, but rather appears to reflect an implementation of two different odor coding strategies. A similar observation regarding shared odor coding in both tracts has recently been reported in ants (Brandstätter and Kleineidam, 2011). This finding is supported by the observation that m-PNs exhibited complex responses to constant odor stimuli with alternating phases of excitation and inhibition, while l-PNs preferably showed stereotypic phasic–tonic responses (Müller et al., 2002). Using the same recording technique, Krofczik

et al. (2009) found that PNs in the l-ACT exhibited higher odor specificity, while m-ACT PNs were more broadly tuned. In addition, they found that l-ACT PNs exhibited suppressive responses to mixtures, that is, the response to the mixture was smaller than the response to the individual components. In contrast, m-ACT PNs showed hypoadditive mixture responses (mixture response resembled the strongest component response). Yamagata et al. (2009) recorded Ca^{2+}-activity of presynaptic PN boutons in one of their target areas, the MB calyx region. They could confirm narrow vs. broad odor tuning as well as suppressive vs. hypoadditive mixture coding in l- and m-ACTs as reported by Krofczik et al. (2009). In addition, they analyzed how the neuronal responses in both pathways depend on odor concentration. They found strong concentration dependence in m-ACT PN responses in the tested concentration range, while l-ACT PN responses showed only weak concentration dependence, but responded already at very low concentrations.

Taken together, l-ACT PNs show narrower odor tuning, little concentration dependence, and suppressive mixture coding, while m-ACT PNs exhibit broader odor tuning, strong concentration dependence, and hypoadditive mixture coding (**Table 1**). Hence, it appears that PNs in the l-ACT are more suited for coding the identity of an odorant but less so for representing its intensity, while PNs in the m-ACT exhibit a complementary coding strategy.

In spite of the detailed morphological and functional descriptions of the l- and m-ACT pathways, it is still unclear how the characteristic coding strategies found therein are brought about. Lateral inhibition and gain control are well-described properties of neuronal processing in the insect AL which seem to be particularly well suited as candidate mechanisms mediating l-ACT-like narrow odor tuning and concentration invariance. Lateral inhibition has been shown to enhance odor discriminability in the insect olfactory system (Linster and Smith, 1997; Stopfer et al., 1997; Perez-Orive et al., 2004; Wilson and Laurent, 2005). Non-uniform lateral inhibition acting between specific glomeruli has been demonstrated to improve odor discrimination in a computational model (Wick et al., 2010). Computational modeling also suggested that in honeybees, the strength of the inhibitory connection between pairs of glomeruli matches the correlation between their response spectra (Linster et al., 2005). Correlation-dependent lateral inhibition has been shown to significantly increase the performance of a machine-learning classifier framework inspired by the honeybee olfactory system (Schmuker and Schneider, 2007).

A second role of inhibition in the AL enables gain control, that is, compression of dynamic range. Gain control can be implemented by global, recurrent inhibition. Several studies have stressed that odorant discrimination can benefit from recurrent feedback inhibition (Stopfer et al., 2003; Olsen and Wilson, 2008; Asahina et al., 2009). Morphological and functional studies showed that in *Drosophila* certain populations of local inhibitory interneurons in the AL connect only specific sets of glomeruli, while other populations exhibit non-specific projections to virtually all glomeruli (Silbering et al., 2008; Chou et al., 2010; Seki et al., 2010), indicating that in *Drosophila*, specific lateral inhibition and global gain control are probably mediated by distinct populations of neurons, potentially belonging to two different functional subsystems. There is much less data available on local interneuron morphology in the honeybee AL, but LNs in the honeybee exhibit manifold morphology and connectivity patterns (Flanagan and Mercer, 1989; Fonta et al., 1993; Sachse and Galizia, 2002). Some LNs have been shown to target exclusively glomeruli in either the l- or m-ACT part of the AL, supporting the idea of partially separated processing of odor information across pathways (Meyer, 2011).

Table 1 | Response properties of boutons in the l-ACT and m-ACT pathways.

l-ACT	m-ACT
Concentration-invariant	Concentration-dependent
Narrow odor tuning	Broad odor tuning
Suppressive mixture responses	Hypoadditive mixture responses

In this contribution, we approach the question of differential coding strategies and potential benefits in a model study. Our specific aim was to analyze whether lateral inhibition and gain control could mediate the different coding strategies that have been observed in the l- and m-ACT pathways in the honeybee brain. To this end, we developed a computational network model that allowed us to implement a complementary odor code in dual pathways, starting from a model that is based on the concept of virtual sensors (Schmuker and Schneider, 2007). The model uses a fixed network scheme and identical neuronal resources but different parameters for the local inhibitory processing. Moreover, the presence of dual olfactory pathways suggests that they provide a substantial evolutionary advantage in the ecological niche which honeybees occupy. The model we present here allowed us to identify specific advantages of dual pathway odor coding for odor-guided navigation and foraging.

MATERIALS AND METHODS

CONCEPT OF THE NETWORK MODEL

We based our network model on a previously published model for processing multidimensional data inspired by the insect olfactory system (Schmuker and Schneider, 2007). The model has no temporal component and hence neglects the fine temporal scale of odor responses. Odor-evoked activity patterns form temporal trajectories in multidimensional space which reach a stationary point shortly after stimulus onset (Stopfer et al., 2003; Galán et al., 2004; Silbering et al., 2008). Activity values in our model reflect the average activity of a neuronal population at this stationary point, and hence capture only the spatial component of olfactory coding across glomeruli (we elaborate on the limitations of this approach in Discussion).

Figure 1A outlines the connectivity of the network model. Olfactory receptor neurons (ORNs) project onto PNs and LNs which are organized in glomeruli. LNs project inhibitory connections to PNs in other glomeruli. PNs project their output to higher brain areas (not part of the model), but their output may undergo global feedback inhibition, that is, gain control.

SURROGATE ORN INPUT PATTERNS

We simulated the virtual responses pattern of ORNs to a large set of odorants as previously described (Schmuker and Schneider,

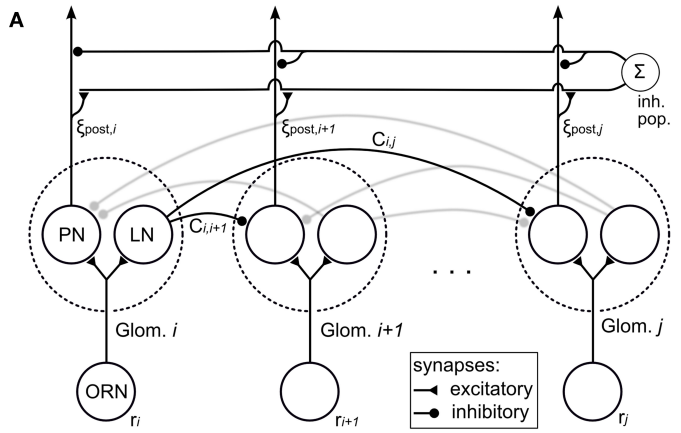

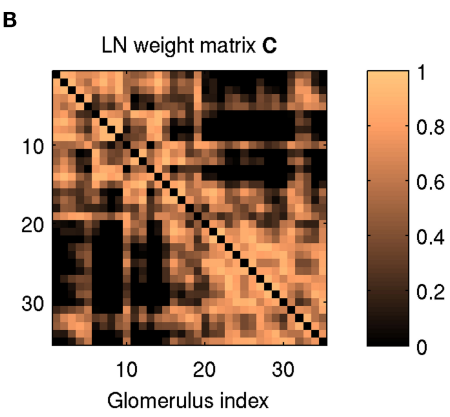

FIGURE 1 | (A) Schematic of the network topology of the model. ORN, olfactory receptor neuron; PN, projection neuron; LN, local (inhibitory) neuron. ORNs project to PNs and LNs in one glomerulus. LNs exert inhibition upon PNs in other glomeruli (Eq. 4). An inhibitory population downstream of the lateral inhibition network takes global input and exerts global inhibitory feedback (see Eqs 5 and 6). Some of the connections are shaded for better overview. **(B)** The weight matrix **C** used for lateral inhibition (cf. Eq. 4).

2007). Using the chemical structure of 836 odorants from the Sigma-Aldrich Flavors and Fragrances Catalog (Sigma-Aldrich, 2004) we calculated a set of 184 physico-chemical descriptors for each individual odorant using the "MOE" software package version 2005.6 (Chemical Computing Group, Montreal). Next, we trained a self-organizing map (SOM) on the 184-dimensional chemical space to mimic the evolution of the olfactory receptor repertoire using the software package SOMMER (Schmuker et al., 2006; http://sommer.sourceforge.net). Each SOM unit was then used as a "virtual receptor" to encode odorants. A virtual receptor is a point $\mathbf{p}$ in data space, and its response r to an odorant $\mathbf{s}$ is determined as a function of the distance between $\mathbf{p}$ and $\mathbf{s}$. The response r_i of the ith virtual receptor – and hence the ith ORN – to an odorant $\mathbf{s}$ is described by Eq. 1

$$r_i = 1 - \frac{d\left(\mathbf{s}, \mathbf{p}_i\right) - d_{\min}}{d_{\max} - d_{\min}}, \tag{1}$$

where $\mathbf{p}_i$ represents the coordinates of the ith receptor, $d(\mathbf{s}, \mathbf{p}_i)$ denotes the city-block distance (i.e., the sum of absolute coordinate differences, Minkowski metric with $k = 1$), $d_{\min}$ and $d_{\max}$ are the minimal and maximal distance between $\mathbf{s}$ and any $\mathbf{p}_i$. Hence, $r_i = 0$ if $d(\mathbf{s}, \mathbf{p}_i)$ is maximal and $r_i = 1$ if $d(\mathbf{s}, \mathbf{p}_i)$ is minimal. The number of virtual receptors can be adjusted by using an SOM with the desired number of units. In this study, we used a toroidal SOM layout with 5×7 units, resulting in 35 virtual receptors/ORNs. The influence of receptor count on stimulus classifiability has been studied previously (Schmuker and Schneider, 2007). Although our model can be efficiently evaluated and supports rapid exploration of large parameter spaces, we limited the model to use only 35 receptors, because increasing the number of receptor types from 35 to 96 yielded only a comparably small improvement in stimulus representation.

CONCENTRATION DEPENDENCE OF PN ACTIVATION

The previous model did not take into account odor concentration, so we needed to extend the model to support odor intensity. In a first step, we account for dynamic range of neuronal activation in the PN layer by expressing the presynaptic activation pattern $\boldsymbol{\xi}$ of PNs by ORNs in terms of a logarithmic transfer function of receptor activation pattern $\mathbf{r}$

$$\boldsymbol{\xi} = \ln\left(\mathbf{r} + 1\right), \tag{2}$$

based on experimental observations in the fruit fly where PN responses saturate with linearly increasing ORN firing (Bhandawat et al., 2007; Kazama and Wilson, 2008). Experimental evidence suggests that short term depressive synapses between ORNs and PNs provide the mechanism that underlies rate compression in the fly (Kazama and Wilson, 2008). Logarithmic transfer functions are also compatible with psychophysical observations as expressed in the Weber–Fechner law.

In a second step, concentration dependence of PN responses was modeled by scaling the presynaptic activation $\boldsymbol{\xi}$ by the decadic logarithm of odor concentration Δ,

$$\boldsymbol{\xi}_{\text{conc}} = \frac{\boldsymbol{\xi}}{1 - \log_{10}\Delta}. \tag{3}$$

We chose to represent odor concentration using this logarithmic description in order to obtain concentration values that are compatible with the physiological study which the present work was inspired from (Yamagata et al., 2009). In that study, odor concentration was reported as a dilution factor of the odorant in solvent. Dilutions ranged from $\Delta = 10^{-5}$ to $\Delta = 10^{0}$, and we adopted these values for the present work. In addition, this kind of scaling is compatible with physiological measurements of concentration-dependent responses in the honeybee AL (Sachse and Galizia, 2003), in which the magnitude of the glomerular response signal scaled with the logarithm of odor concentration.

LATERAL INHIBITION

Many studies have shown that lateral inhibition in the AL plays an important role in olfactory information processing (see, e.g.,

Stopfer et al., 1997; Perez-Orive et al., 2004; Wilson and Laurent, 2005). In this study, we used correlation-dependent lateral inhibition, as suggested by Linster et al. (2005) and in agreement with the previous incarnation of the model (Schmuker and Schneider, 2007).

We implemented lateral inhibition by subtracting from the activation of each PN the summed activation of all other PNs, weighted by correlation between the ORN activations over all odorants in the data set. Accordingly, we obtain the output activation pattern $\boldsymbol{\xi}_{\text{post}}$ from the concentration-dependent input activation pattern $\boldsymbol{\xi}_{\text{conc}}$ as

$$\boldsymbol{\xi}_{\text{post}} = \boldsymbol{\xi}_{\text{conc}} - q \cdot \frac{\mathbf{C} \cdot \boldsymbol{\xi}_{\text{conc}}}{n}, \tag{4}$$

where n is the number of virtual receptors, $\mathbf{C}$ is the correlation matrix with $C_{i,j}$ containing the Pearson correlation coefficient for the responses of the ith and jth ORN, and q is a scaling factor which allows to adjust the overall strength of inhibition. Elements on the diagonal as well as negative entries of $\mathbf{C}$ were set to zero to avoid self-inhibition and to reflect the fact that inhibited inhibitory interneurons have no post-synaptic effect. **Figure 1B** depicts the weight matrix $\mathbf{C}$. It is calculated on the basis of receptor responses to all single odorants in the data set and thus reflects the similarity structure in the input space. It is constant throughout this study. Note that the structure of lateral inhibition is controlled by $\mathbf{C}$, but q is factored out to allow control of the overall strength of inhibition. The factor n was introduced in the previous work to enable comparison of q values over different receptor counts. Although n is constant in this study, we kept it as a parameter for consistency with previous work.

FEEDBACK INHIBITION

Our implementation of feedback inhibition was motivated by two experimental observations. First, it has been shown that l-ACT boutons respond to very weak odor concentrations, while m-ACT PN boutons do not (Yamagata et al., 2009), meaning that PNs in the l-ACT react more sensitive to weak stimuli compared to m-ACT PNs. Second, in the honeybee, the *input* to l-ACT PNs in the AL exhibits strong dependence on odor concentration. This has been demonstrated in a study which measured the Ca^{2+} signal in backfilled PNs in the AL, where dendritic Ca^{2+} activity likely constituted the majority of the signal (Sachse and Galizia, 2003). It hence follows that gain control in the honeybee AL must act downstream of PN dendrites, but upstream of PN boutons.

In order to capture this behavior, we extended the model with a two-step gain control mechanism consisting of a sensitivity boost and subsequent feedback inhibition. Our implementation of gain control is described as

$$\boldsymbol{\xi}_{\text{out}} = \frac{\boldsymbol{\xi}_{\text{post}} \cdot \beta}{\rho}, \tag{5}$$

where β denotes the sensitivity boost factor, and $\rho \geq 1$ describes the activity of a population of inhibitory neurons that mediates global feedback inhibition on PN output (see **Figure 1A**) in a divisive manner. The inhibitory population receives input from all PNs; its activity is described as

$$\rho = \begin{cases} 1, & \text{if } \left|\boldsymbol{\xi}_{\text{post}}\right| \leqslant \theta \\ \frac{\left|\boldsymbol{\xi}_{\text{post}}\right|}{\theta}, & \text{if } \left|\boldsymbol{\xi}_{\text{post}}\right| > \theta \end{cases}. \tag{6}$$

$|\boldsymbol{\xi}_{\text{post}}|$ denotes the L1-norm of the PN output pattern (the sum of the activity of all PNs), θ the (non-zero) activation threshold of the inhibitory neurons. The case of $\rho \equiv 1$ describes the case where no feedback inhibition is present in the network. The inhibitory population responds as total PN output exceeds θ, which we set to the average level of PN output for all patterns in our data set without lateral inhibition and gain control. We normalized the amount of inhibition by θ, so that $\rho \approx 1$ when $|\boldsymbol{\xi}_{\text{post}}| \approx \theta$. Taken together, gain control in our model globally boosts the PN signals if summed activity is weak and attenuates them if it the summed response is strong.

In our physiological observations, the output from m-ACT neurons at the highest concentration was in the same range as the response from l-ACT neurons over concentrations (Yamagata et al., 2009). Hence, we set β to 6 throughout this study, compensating the concentration discount for the lowest concentration (10^{-5}) in Eq. 3. Thus, when gain control is effective, the amplitude of an input pattern at the lowest concentration matches the amplitude of the same pattern at the highest concentration without gain control.

RESULTS

Stimulus representation differs in a characteristic way across two parallel pathways in the honeybee olfactory system (**Table 1**; Yamagata et al., 2009). Our aim was to identify network properties underlying the observed differential odor coding, and to investigate the potential benefit of such a dual coding strategy. To this end, we used a computational network inspired by the neural circuits in the olfactory system of the honeybee, based on a model of olfactory processing in the AL (Schmuker and Schneider, 2007). We investigated the roles which glomerulus-specific lateral inhibition and non-specific feedback inhibition ("gain control") may play in creating the distinct coding properties in the l- and m-ACT. Since comprehensive characterizations of the molecular receptive fields of the honeybee's ORNs are not available, we used surrogate data as input to the AL (see Materials and Methods for details).

EFFECT OF LATERAL INHIBITION STRENGTH AND GAIN CONTROL ON THE WIDTH OF ODOR TUNING CURVES

It has long been known that lateral inhibition effects contrast enhancement and supports a sparse stimulus code in sensory processing of various modalities. In the honeybee, PN boutons in the l-ACT pathway respond to fewer odorants and thus exhibit narrower odor tuning profiles than their m-ACT counterparts – in other words, l-ACT PNs exhibit a sparser odor tuning code (Yamagata et al., 2009). In our model, increasing lateral inhibition in the absence of gain control leads to an overall reduction in activity, likely making it more difficult for downstream neurons to detect the response patterns (**Figure 2A**, upper row). Obviously, gain

control (Eqs. 5 and 6) can act as a mechanism to recover the activity pattern (**Figure 2A**, lower row). With gain control enabled, the activity pattern retains its strength, even when lateral inhibition is strong ($q > 1.0$). At the same time, the sparsening effect of lateral inhibition can be observed: many PNs respond when the lateral inhibition is weak (q is low), but when q increases the activity pattern gets sparser and its contrast is enhanced.

We then tested whether lateral inhibition as described by Eq. 4 can account for narrow odor response profiles and sparse odor coding in l-ACT PNs. **Figure 2B** illustrates the tuning profile of a single PN with respect to the entire set of 836 odors in our database. Without gain control, the odorant response profiles show overall reduced response magnitude, but their shapes are not changed (**Figure 2B**, left panel). But, with gain control in effect, odorant tuning curves become narrower as lateral inhibition is increased (**Figure 2B**, right panel). Thus, we could achieve substantial narrowing of the odorant response profiles only when gain control was in effect and, at the same time, lateral inhibition was strong.

One might argue whether increasing input gain alone may also be sufficient to produce narrower tuning curves in presence of strong lateral inhibition. We tested this hypothesis by increasing sensitivity alone without applying feedback inhibition ($\beta = 6$ and $\rho = 1$ in Eq. 6). The resulting tuning curves did not become narrower, only their amplitude was increased (**Figure 2C**). This observation further supports that both, increased input sensitivity *and* gain control by feedback inhibition are required to reproduce narrow odor tuning when lateral inhibition is strong.

Lateral inhibition affects the discriminability of odorant representations in PNs. We used the Euclidean distance between PN response patterns as an estimate for their discriminability. Distance was calculated for each pair of odors, so that we ended up with 836*835*0.5 = 349,030 distance values for each value of lateral inhibition strength q. Two observations can be made from the distribution of distances (**Figure 3**). First, when gain control was in effect, increasing the strength of lateral inhibition also caused the distance between odorant representations to increase. Second, when there was no gain control, the distance between odorant representations slightly decreased with increasing lateral inhibition. This observation is a consequence of the decreasing overall amplitude of response patterns with increasing lateral inhibition. This result indicates that both mechanisms, gain control *and* lateral inhibition, need to be combined to achieve improved pattern separability, and hence discriminability by downstream neurons.

This interpretation relies on the assumption that neuronal responses are not invariant to scaling, that is, the assumption that small differences in the signal are more difficult to discern than large differences. Assuming that stochastic fluctuations in the stimulus, the transduction chain, synaptic release, external conditions, etc., do not or only weakly scale with signal amplitude, they will affect weak signals more than strong ones. In consequence, large differences in activity between PNs will be more reliably detected

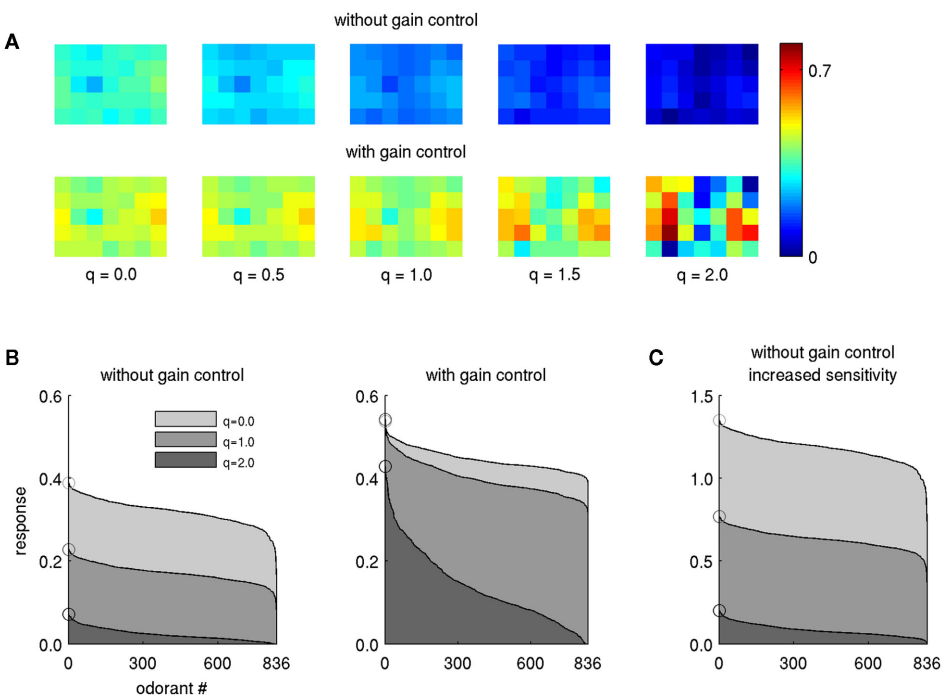

FIGURE 2 | Effect of lateral inhibition on PN response profiles. (A) The pattern evoked by a single odor (butyl proprionate, modeled concentration 10^{-1}) transformed by lateral inhibition of different strengths, indicated by the factor q. Each square in the 5 × 7 grid corresponds to the PN response within one model glomerulus. **(B)** Response magnitude in PN no. 17 (fourth column, second row) for all 836 odors for three levels of lateral inhibition, without (left) and with gain control and increased sensitivity ($\beta = 6$, right). The odors are arranged along the abscissa with descending response strength such that the strongest odorant is always displayed on the left. Circles denote the magnitude of the leftmost odorant. **(C)** Same as **(B)** without gain control, but with PN sensitivity increased by the same factor ($\beta = 6$). Note the different scale of the ordinate.

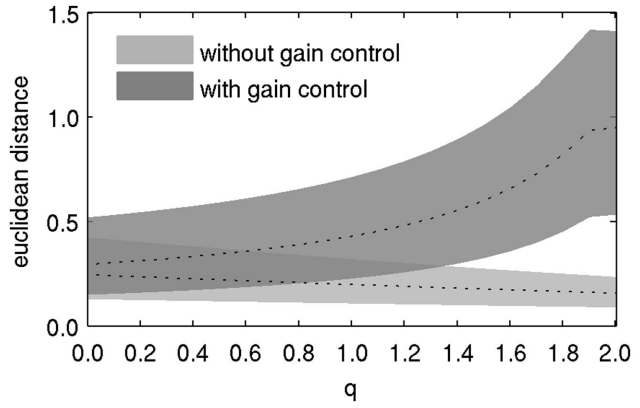

FIGURE 3 | Discriminability of odor representations in PNs, as indicated by the average pairwise euclidean distance between PN response patterns to all odorants. Dotted lines denote the median, shaded areas denote the difference between the 10th and the 90th percentile of the distribution of distances.

than weak differences. To demonstrate this aspect would require a more elaborate model which is outside the scope of this study.

Taken together, lateral inhibition alone leads to reduced responses without substantial sharpening of PN tuning curves. It only improves odor discrimination together with gain control in the scope of our model.

REPRODUCTION OF PHYSIOLOGICAL OBSERVATIONS

We wondered whether we could also reproduce the characteristic differences in odor concentration and mixture representation across l- and m-ACT that have been described in physiological studies (**Table 1**). In the following, we describe the extent to which our model is able to reproduce physiological observations regarding representation of odor concentration and mixtures of odors. Since we implemented gain control and lateral inhibition as separate mechanisms, we were able to gauge their effects on odor coding separately.

In our previous physiological experiments, we observed positive concentration dependence in m-ACT PNs and no or even negative concentration dependence in l-ACT PNs (Yamagata et al., 2009). We first wanted to investigate to what extent gain control alone could reproduce our experimental observations on concentration and mixture coding, when there is no lateral inhibition involved ($q = 0$ in Eq. 4). Without gain control, the PN responses naturally exhibited strong concentration dependence, matching the behavior of m-ACT PNs (upper row in **Figure 4A**). With gain control, this dependence was virtually eliminated (lower row in **Figure 4A**). In many cases we even observed negative concentration dependence, as for example in the upper trace in **Figure 4B**. Negative concentration dependence occurs when a PN is strongly activated by its input, but gets attenuated through feedback inhibition as the other PNs increase their responses with rising concentration. It was weak in magnitude in the cases where it was present. Negative concentration dependence was also observed physiologically in a large fraction of l-ACT PN boutons (cf. Figure 4 in Yamagata et al., 2009).

In order to depict the general effect of feedback inhibition on intensity representation, we followed the approach of our physiological study by estimating the representation of different odor concentrations in individual PNs. To this end, we performed linear regression of the concentration response curves (**Figure 4B**) for the response of each of the 35 PNs to each of the 836 odors, obtaining $836 \cdot 35 = 29{,}260$ slopes. We then compared the distribution of slopes for the two conditions with and without gain control (**Figure 4C**). Without gain control the regression lines had strictly positive slopes, but when gain control was enabled, the regression slopes took small values close to zero, both in the positive and negative range. Hence, the strong concentration dependence observed in m-ACT PNs was reproduced when gain control was disabled, while l-ACT-like concentration coding could be achieved with gain control through feedback inhibition.

Responses of l-ACT PNs to mixtures of two odorants have been described to be weaker than their responses to each of the components (suppressive mixture coding), while in m-ACT PNs mixture responses have been observed to be as strong or stronger than the response to each component (hypoadditive mixture coding; Krofczik et al., 2009; Yamagata et al., 2009). We systematically varied the strength of lateral inhibition in absence and in presence of gain control to investigate the role of those network properties in reproducing characteristic differences in mixture coding across pathways. To this end, we implemented mixtures of two odorants as linear superposition of their patterns (Eq. 7)

$$\boldsymbol{\xi}_{\text{mix}} = \ln\left(\mathbf{r}_{\text{A}} + \mathbf{r}_{B} + 1\right), \tag{7}$$

with $\mathbf{r}_{\text{A}}$ and $\mathbf{r}_{\text{B}}$ the receptor patterns of the two components A and B, and $\boldsymbol{\xi}_{\text{mix}}$ the activity pattern in response to the mixture. As in Eq. 2, we used a logarithmic transfer function to account for the dynamic range of neuronal activation. Activity patterns are combined before applying the logarithmic transfer function which models PN activation. In the biological scope, this approach corresponds to linearly adding the responses on the receptor level, which is in line with physiological findings in various organisms. For example, Tabor et al. (2004) observed in zebrafish that activity patterns evoked by odor mixtures in afferents to the olfactory bulb could be predicted from the component patterns, suggesting that mixture interactions in the peripheral olfactory system exhibit only weak non-linearity. Similar findings have been described in *Drosophila* (Silbering and Galizia, 2007), the moth *Spodoptera litoralis* (Carlsson et al., 2007), and the honeybee (Deisig et al., 2006). To illustrate the effect of mixing odorants, we depict response patterns of two odorants (acetaldehyde and butyl proprionate) and their mixture without gain control in **Figure 5A**. We used the (virtual) concentration of 10^{-1}, the same concentration that we used in the mixture experiments in our experimental study (Yamagata et al., 2009). PNs exhibited hypoadditive mixture coding, as the response of each PN to the mixture was higher than to each of the components, although the mixture response is not a strictly additive superposition of both components because of the PN's logarithmic transfer function (Eq. 7).

We then checked the influence of gain control and lateral inhibition on mixture coding. **Figure 5B** shows the response patterns

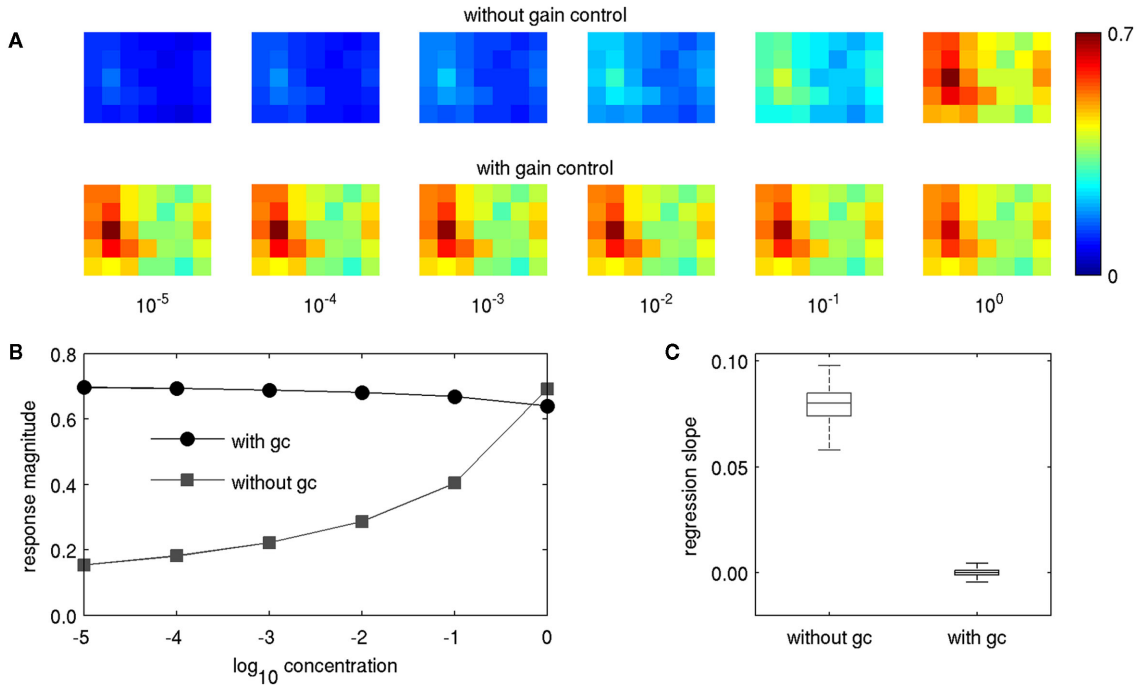

FIGURE 4 | Concentration dependence of PN response patterns with and without gain control. (A) The pattern for butyl laurate for concentrations of 10^{-5}–10^{0} without gain control (upper row) and with gain control (lower row). **(B)** Response magnitude of the PN in glomerulus no. 8 (the strongest responding glomerulus, second column, third row) with and without gain control. **(C)** Box plot of regression slopes of concentration response curves as in **(B)**, but pooled for the entire data set, over all odors, and PNs. Lines in boxes denote upper quartile, median, and lower quartile, "whiskers" denote 10th and 90th percentile.

for components and mixtures with gain control, and for three different levels of lateral inhibition (controlled by the factor q). Gain control generally promoted suppressive mixture coding. Lateral inhibition led to increased contrast of the components as well as the mixtures (cf. **Figure 2**), but it was difficult to discern an effect on mixture coding for $q = 1$. A clear effect became visible for $q = 2$, but it is difficult to estimate from the visualization of the response pattern alone whether this condition supported suppression or hypoadditivity. A significant portion of PNs were silenced by strong lateral inhibition for $q = 2$, resulting in sparse activity in the component-evoked pattern, and hence lower overall gain. With lower total gain, gain control became less effective (see Eq. 6), and its suppressive influence on mixture representation became smaller. This observation indicates that values of $q \geq 2$ represent an extreme case and are most likely not useful in reproducing experimental observations.

We wanted to obtain a quantitative assessment of whether PNs express suppressive or hypoadditive responses to a mixture of two components A and B over a range of parameters. To this end, we calculated for each PN the index for mixture additivity κ (Krofczik et al., 2009; Eq. 8)

$$\kappa = \frac{\xi_{i,\,\mathrm{mix}} - \max\left(\xi_{i,\,\mathrm{comp\,A}}, \xi_{i,\,\mathrm{comp\,B}}\right)}{\xi_{i,\,\mathrm{mix}} + \max\left(\xi_{i,\,\mathrm{comp\,A}}, \xi_{i,\,\mathrm{comp\,B}}\right)}, \tag{8}$$

with $\xi_{i,\,\mathrm{comp\,A}}$ and $\xi_{i,\,\mathrm{comp\,B}}$ the response of the PN with the index i to the component A and B, respectively, and $\xi_{i,\,\mathrm{mix}}$ the response of PN i to the mixture. Hence, values of $\kappa > 0$ indicate hypoadditive mixture representation, while values of $\kappa < 0$ indicate suppressive mixture representation. **Figure 5C** shows the distribution of κ across PNs for the mixture from **Figure 5A**, with and without gain control, for q in the range between zero and two. Gain control supported suppressive mixture representation ($\kappa < 0$) over a large range of q. Although the values are still close to zero, the median as well as the 10th and 90th percentiles of the distribution of κ were below zero, indicating a clear trend to suppressive responses. In contrast, mixtures were represented in a hypoadditive way ($\kappa > 0$) when there was no gain control. Strikingly, the amount of lateral inhibition had virtually no effect on the mixture responses being hypoadditive or suppressive over a large range of inhibition strength. That stereotypic behavior began to break down only as q approached the extreme value of 2, when no clear trend toward hypoadditive or suppressive mixture coding could be discerned anymore, confirming our qualitative observations from **Figure 5B**. Hence, the presence of gain control appears to be the dominant factor mediating the characteristics of mixture coding over a large range of parameter values.

Natural odors are typically blends of many different components (Knudsen et al., 2006; see Raguso, 2008 for a review), and the ability to discriminate blends is important for foraging (Wright et al., 2002). Honeybees can distinguish fine differences in blend composition foraging (Wright et al., 2005), and it has shown that AL processing plays an important role in mixture discrimination (Deisig et al., 2010). As lateral inhibition and gain control enhance

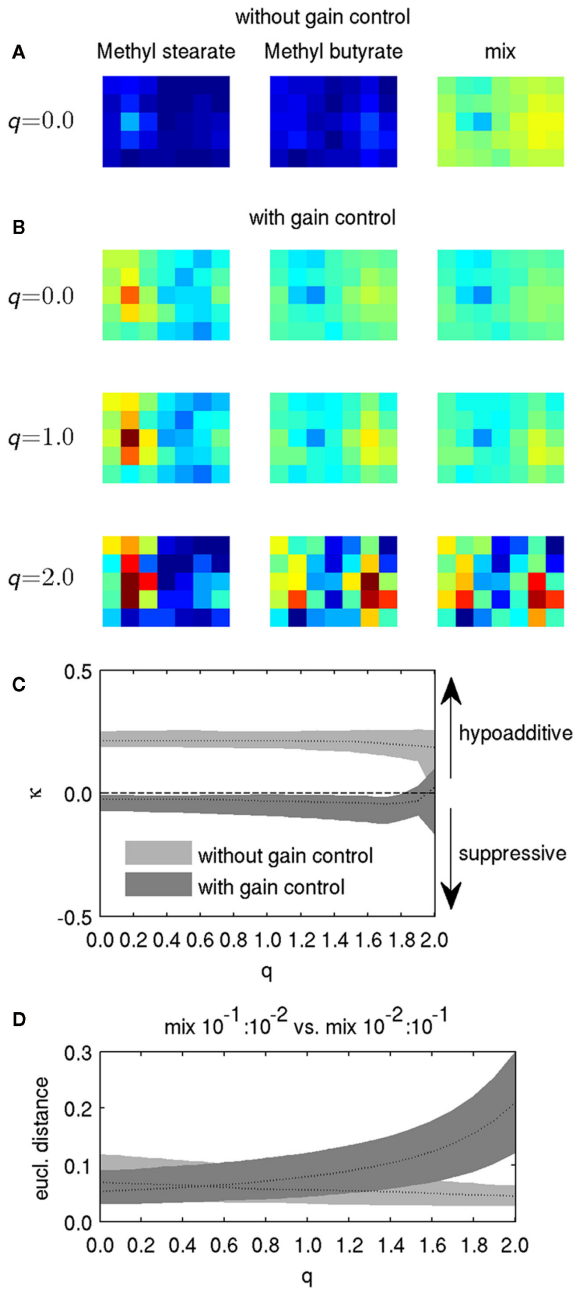

FIGURE 5 | Effect of lateral inhibition and gain control on mixture representation. (A) Response patterns without gain control and $q = 0$ for two odorants at concentration 10^{-1} and their mixture. **(B)** Response patterns for the odorants and the mixture from **(A)**, but with gain control and for different values of q. Color scaling in **(A,B)** is identical to **Figures 2A** and **4A**. **(C)** The mixture index κ for different values of q, with and without gain control, calculated for all glomeruli for the mixture of Acetaldehyde and Butyl propionate. Black dotted lines indicate the median for all glomeruli, for different values of q. Shaded areas denote the 10th and the 90th percentile of κ. The dashed line indicates $\kappa = 0$, where mixtures of odorants evoke exactly the same response as the strongest components, and hence exhibit neither hypoadditive nor suppressive responses. **(D)** Euclidian distance between binary mixtures with slightly different component concentrations. Dotted line: Median of 100 randomly picked odorant pairs, shaded areas denote 10th and 90th percentile.

pattern discriminability (cf. **Figure 3**), we hypothesized that these mechanisms also affect mixture discrimination. We tested whether our model reproduces this observation by calculating the distance between artificial binary blends with slight differences in the proportion of components. To this end, we randomly picked odorant pairs from our data set and created two binary mixtures from each pair that differed slightly in their component ratio. In one mixture the ratio of concentration between odor A and odor B was $A:B = 10^{-1}:10^{-2}$, while the ratio was reversed in the other mixture ($A:B = 10^{-2}:10^{-1}$). We calculated the Euclidean distance between the response patterns evoked by the two mixtures after processing them with various levels of lateral inhibition. As with single odors, discriminability of mixtures was enhanced for high levels of lateral inhibition in combination with gain control (**Figure 5D**), thus favoring the discrimination of weak differences in blends of odorants.

Taken together, lateral inhibition enhanced discriminability of mixtures on the level of PN response patterns, but it hardly affected whether mixtures were represented in an additive or subtractive way at single PNs. The latter aspect of mixture coding was mediated mainly by the presence or absence of gain control, which is required to avoid that the signal gets lost when lateral inhibition is strong.

MODEL-BASED EXPLANATION OF AN ATYPICAL EXPERIMENTAL OBSERVATION

The mechanism we suggest here helps to understand an important detail in our previous experimental observations on PN boutons in the MB. There was one odor, hexanal, which exhibited concentration dependence more in l-ACT PN boutons than in m-ACT boutons, and was more sparsely represented in m-ACT PN boutons than in l-ACT PN boutons, hence exhibiting "inversed" behavior (cf. Figure S3 in Yamagata et al., 2009). Both effects can be explained in a straightforward manner from our model, considering the particular response pattern that hexanal evokes in the AL. Hexanal activates comparably few glomeruli in the part of the l-ACT that is accessible to imaging (Sachse et al., 1999). In the m-ACT, sparse input will lead to sparse output, but in the l-ACT, PNs in glomeruli which receive only weak excitatory input may exhibit stronger responses at their output through the sensitivity boost that gain control implies (factor β in Eq. 5), apparently reducing sparseness compared to m-ACT PNs. This explains why hexanal is sparsely represented in m-ACT PN boutons, but exhibits less sparse representation in l-ACT PNs.

Along a similar line of thought, our model explains the other effect we observed, namely concentration-dependent representation of hexanal in otherwise concentration-insensitive l-ACT boutons. In the l-ACT, gain control will not exert much influence when only few PNs are activated, as the strength of gain control depends on the sum of activation in all PNs (Eq. 6). So gain control is less effective for odors which activate few PNs (like hexanal), and in consequence those odors will exhibit concentration dependence also with l-ACT-like network parameters.

DISCUSSION

In the present study, we have analyzed putative network mechanisms which may generate the characteristic features of odor

Table 2 | Network properties reproducing l-ACT and m-ACT-like coding characteristics and their functional significance.

l-ACT	m-ACT
Strong lateral inhibition good odor discrimination	Weak lateral inhibition weak odor discrimination
Strong gain control weak concentration discrimination	Weak gain control good concentration discrimination

coding that have been described across two parallel pathways in the honeybee olfactory system. We have demonstrated that the complementary features of odor coding at the level of the MB input found across the two pathways (**Table 1**) can be captured in a single model network by using different parameter settings (**Table 2**). We could reproduce coding properties observed in PNs of the l-ACT, namely narrow odor tuning, concentration invariance, and suppressive mixture coding by strong lateral inhibition and strong gain control in the model network. On the contrary, weak lateral inhibition and absence of gain control reproduced coding properties of PNs in the m-ACT, namely broad odor tuning, concentration-dependent responses, and hypoadditive mixture representation. As a result in the biological scope, this finding is compatible with the assumption that in the honeybee AL, the l-ACT/m-ACT pathways represent two partly segregated subnetworks where parameters are tuned to differential functional properties using the same neuronal ground plan. This constrains current hypotheses on the connectivity within and between both pathways (see Galizia and Rössler, 2010 for a review). Our results also suggest dual pathway odor coding may be of particular value in odor-guided foraging, when precise information about stimulus intensity and identity is required simultaneously.

Several assumptions underlying our model are vital for the validity of these interpretations. For the present study we assumed that odor selectivity is similar for ORNs that innervate glomeruli of the l-ACT and those innervating m-ACT. Also, we assumed homogenous odor sensitivity in ORNs that innervate glomeruli of either tract. In our model, tract-specific differences in odor sensitivity arise only at the level of PNs and result from computations within the AL network and (potentially) feedback from downstream neuron populations (see below). To our best knowledge, there is currently no experimental evidence that contradicts any of these assumptions.

CANDIDATE NEURON POPULATIONS IN THE HONEYBEE BRAIN MEDIATING FEEDBACK INHIBITION ON THE AL OUTPUT

Our modeling results suggest that gain control is important to explain the coding differences between l-ACT and m-ACT. In order to test this hypothesis in future studies, it is important to know which neuron populations could mediate the gain-controlling feedback inhibition in the honeybee olfactory pathway. Evidence from morphological and physiological studies in the honeybee brain suggests three possibilities.

The first possibility is that gain control is provided by inhibitory feedback neurons in the protocerebro-calycal tract (PCT). PCT neurons have post-synaptic profiles in the alpha-lobe of the MB and project up to the calyx region (Grünewald, 1999; Okada et al., 2007). In the MB calyx, PCT neurons form reciprocal connections to PN boutons within so-called microglomeruli (Ganeshina and Menzel, 2001). At least a part of the PCT neurons are GABAergic. They could mediate the observed feedback inhibition since the physiological observations which we partly ground our model upon are based on measurements of presynaptic PN boutons within these microglomeruli (Yamagata et al., 2009). Dendrites of PCT neurons span large areas in the alpha-lobe, where they can pick up the population response from a large fraction of Kenyon cells (KCs), which in turn may reflect the overall input of PNs to the MB. Hence, the connection PN–KC–PCT–PN may form an inhibitory feedback loop. Recently, a similar mechanisms has been proposed in locust, where negative feedback from the MB output back to its input is mediated by a large inhibitory neuron (Papadopoulou et al., 2011), although it is unclear whether that neuron targets only KC dendrites or if it also synapses onto PN axonal boutons.

The second possibility is that feedback inhibition from PCT neurons is mediated directly through the widespread axo-dendritic ramifications that these neurons form in the calyx region. Within microglomeruli, PCT neurons can pick up the excitatory input directly from PNs and feed it back on the boutons through their inhibitory synapses in the entire calyx region (Ganeshina and Menzel, 2001; Hourcade et al., 2010). However, it is unclear how signals traveling on the axo-dendritic structures between microglomeruli would interfere with action potentials arriving from the alpha-lobe.

As a third possibility, gain control could be exerted by local inhibitory interneurons in the AL targeting the PN dendrites very close to the integrating segment, that is, in the core region of glomeruli (Fonta et al., 1993). The dendrites would still be showing concentration-dependent responses, but due to hyperpolarization close to the integrating segment, fewer action potentials might be generated. Hence, Ca^{2+}-imaging of backfilled PNs in the AL would still show concentration-dependent responses, in line with experimental observations (Sachse and Galizia, 2003). At the same time, this arrangement would be compatible with the absence of concentration dependence in l-ACT PN output boutons in the MB that we observed (Yamagata et al., 2009).

The question whether gain control in l-ACT PNs is achieved already in the AL or only at the level of the PN boutons may be resolved by measuring the activity in axons of l-ACT PNs, for example by measuring electrophysiological signals extracellularly in the tract (Brill et al., 2011). Also, more detailed physiological and morphological characterizations of local interneurons in the honeybee AL are required (Meyer, 2011; Meyer and Galizia, 2011).

POTENTIAL RELEVANCE OF DIFFERENTIAL ODOR CODING IN ODOR-GUIDED NAVIGATION AND FORAGING

In our model, lateral inhibition and gain control together can enhance discriminability of single odors (**Figure 3**) and blends (**Figure 5D**). On the other hand, gain control inevitably leads to a loss of intensity information (**Figure 4**). Hence, there exists a trade-off between discriminability of odors and representation of their intensity. The co-existence of a concentration-invariant and highly discriminative coding scheme in the l-ACT and a concentration-sensitive, less discriminant coding scheme in the m-ACT might

indicate an evolutionary adaptation of the honeybee olfactory system to achieve both, concentration sensitivity *and* accurate detection of odors even at low concentration.

Both capabilities are of vital importance for odor-guided foraging. For example, a foraging bee may encounter a faint odor trace emanating from a distant flower. Under these conditions, detection and identification of the odorant requires high sensitivity at low odor concentrations and narrow odor tuning, as found in l-ACT PNs. If the odor is attractive, the bee will try to approach the odor source. Information about the distance to the source is encoded in odor plumes in two factors: odor intermittency, that is, the frequency at which the odor is encountered, and odor concentration (see Riffell et al., 2008 for a review). Although it has been shown that odor intermittency alone can be used to successfully approach the source of a pheromone via "infotaxis" (Vickers, 2000; Vergassola et al., 2007; Moraud and Martinez, 2010), it is clear that the distance to the source of an odor is also encoded in its concentration within plume filaments (Murlis, 1997; Thistle et al., 2004; Zollner et al., 2004) and hence odor concentration is a powerful cue for distance to the source. Information about concentration is encoded in m-ACT PNs. Thus, by integrating the information about odor identity from l-ACT PNs and concentration from m-ACT PNs, the honeybee may be able to detect and approach an odor source more efficiently than when relying on infotaxis alone. In addition, odor concentration can be very high in close vicinity to the odor source, e.g., at the blossom of a flower. Odor discrimination (and hence odor learning) in this scenario requires that the PN response is not saturated. Gain control may counteract saturation of l-ACT PNs, enabling reliable, and concentration-invariant encoding of odor identity even at high concentrations. Interestingly, previous model studies have shown that, indeed, gain control can improve learning of odor identity (Huerta et al., 2004; Nowotny et al., 2005).

If our hypothesis is correct we may predict that a honeybee whose l-ACT is dysfunctional will still be able to approach odor sources, but it should exhibit severe impairments in general odor discrimination, in odor learning at high concentrations, and in odor detection at low concentrations. On the other hand, a honeybee lacking a functional m-ACT should exhibit little impairment in odor detection and discrimination, but it may have an impaired ability to approach an odor source, unless the lack of intensity information conveyed by m-ACT PNs is compensated by another mechanism or other neurons, like for example ml-ACT PNs or multi glomerular PNs (see Galizia and Rössler, 2010 for detailed overview on PN types in the honeybee AL).

The latter point leads to the question why an entire odor coding system like the m-ACT should be used to represent intensity if in principle only one channel would be necessary. However, in an odor rich environment, where one odor needs to be identified out of many odors, and its concentration to be determined, it may be important to have not only a plain intensity signal, but also substantial odor information associated with it. This information may help to bind the concentration information relayed by the m-ACT to the odor identity information relayed by the l-ACT in higher brain regions where both pathways converge. Additional experimental and theoretical studies are needed to test this hypothesis.

RELATION TO ANOTHER PROPOSED MECHANISM FOR CODING ODOR IDENTITY VS. INTENSITY

Stopfer et al. (2003) found that in locust, odor intensity, and identity are represented in parallel within the same spike train, encoded in temporal patterns which change dynamically during odor presentation. Locusts do not have dual olfactory pathways like the l- and m-ACT in honeybees, so the challenge of encoding identity and intensity at the same time may be solved differently in these animals. Nevertheless, the presence of parallel pathways for intensity and identity in the honeybee raises the question of the particular benefit of that dual system. One potential advantage of having a parallel representation of those stimulus features as opposed to encoding it in temporal features is that in a parallel system, identity and intensity can be provided almost immediately after stimulus onset. If those features have to be decoded from a spatio-temporal pattern, a certain amount of time is required until that pattern has evolved to encode a sufficient amount of information. The olfactory system of honeybees signals the presence of behaviorally relevant odors with very little delay in conditioning experiments (Strube-Bloss et al., 2011). In those experiments, MB-extrinsic neurons which encode the value of the stimulus (rewarded vs. non-rewarded) exhibited an odor-driven population response 60 ms after odor onset, which is approximately 20 ms after a specific representation of the odor has built up in the PN ensemble in the AL (Krofczik et al., 2009). Moreover, the representation of the rewarded odor at the MB output became significantly different from unrewarded stimuli as quickly as 140 ms after odor onset. These findings suggest that the olfactory system of the honeybee can indeed benefit from a rapid encoding of odors in the input, and hence from parallel pathways. Moreover, in a natural environment odor plumes are encountered intermittently and concentration is fluctuating very fast when crossing a filament (Murlis, 1997; Thistle et al., 2004; Zollner et al., 2004; Riffell et al., 2008). If odorant concentration and identity are to be linked, it is hence important to measure the maximum odor concentration with high temporal precision. This capability may particularly facilitate odor-guided navigation in an odor rich environment where the animal encounters plumes from different odor sources and needs to distinguish between them.

LIMITATIONS OF THE PROPOSED MODEL IN THE BIOLOGICAL SCOPE

The network model presented in this study considers only the spatial component of the olfactory code and ignores detailed response dynamics. The main reason for this restriction is that our model is partly based on experimental results obtained from Ca^{2+} imaging of PN boutons, which provides only poor estimates of the temporal dynamics of the neural code. However, there is much evidence for the relevance of temporal aspects of the neural code. Stopfer et al. (1997) have demonstrated the importance of odor-evoked oscillations in neuronal activity in the AL for odor discrimination in the honeybee. Krofczik et al. (2009) suggested that PNs may employ a latency code for representing odors, where the temporal delay between stimulus onset and PN response is characteristic for each pair of odor and PN. The temporal structure of odor representation in physiological recordings from AL neurons in locust and fruit fly has been demonstrated to vary strongly with odor identity, even under constant stimulus conditions (Müller et al.,

2002; Stopfer et al., 2003; Wilson et al., 2004). Likewise, intensity has been shown to affect the temporal dynamics of PN responses in a characteristic way (Stopfer et al., 2003; Ito et al., 2009). A large part of the temporal complexity seems to be generated already at the level of ORNs (Raman et al., 2010). A recent modeling study has shown how downstream neurons can extract temporal relationships between the firing of populations in the locust AL (Assisi et al., 2011).

Most studies that deal with olfactory coding in the AL (including the present one) have focused on either the spatial or the temporal aspect of the spatio-temporal code while putting little emphasis on the other. This can be partly explained by the distinct experimental approaches that provide the empirical basis for model approaches. Imaging techniques monitor a spatial arrangement, often with low temporal resolution, while electrophysiological recordings of single neurons or few units provide high temporal resolution but poor spatial information. The existence of a spatial code in the honeybee has been repeatedly proven in glomerular space and is reproducible across animals (Galizia et al., 1999; Wang et al., 2003). However, glomerular response patterns have mostly been obtained by measuring Ca^{2+}-activity, with temporal resolution typically in the range of 5 Hz. This resolution is not sufficient to identify temporal variations in odor representation, which have been described to occur at frequencies starting at 20 Hz using electrophysiological methods (e.g., Stopfer et al., 1997; Laurent et al., 1998; Wilson and Laurent, 2005). On the other hand, electrophysiological studies either have very small spatial coverage, or lack the possibility to determine the exact spatial location of the signal source, e.g., in the case of local field potentials. Hence, a direct comparison of how spatial (that is, glomerular) and temporal coding interact in the insect olfactory system is still missing. Given the current experimental evidence, it is difficult to reconcile the debate whether the spatial or the temporal component of odor coding is more important in odor discrimination. It is also likely that different insect species have evolved different odor coding strategies, given the considerable difference in the anatomical layout of their olfactory systems (e.g., number of glomeruli, functional separation between structural glomeruli, single vs. dual pathway systems, multi- vs. uniglomerular PNs, see Martin et al. (2011) for a review), possibly with different emphasis on spatial and temporal encoding of odors.

In this context, it is interesting to note that in the honeybee, temporally complex responses are observed frequently in m-ACT PNs (Müller et al., 2002), but rarely in l-ACT PNs (Krofczik et al., 2009). This observation may indicate that l-ACT and m-ACT are not only different in the spatial component of odor representation, but also in the temporal component. Moreover, the m-ACT is shared between locust, *Drosophila*, and honeybee, while the l-ACT is unique to *Hymenoptera*, notably bees, and ants (Galizia and Rössler, 2010). The difference in temporal aspects of the neural code between l- and m-ACT will be a fascinating topic for future studies as simultaneous observations from PNs in both pathways with suitable temporal resolution become available (Brill et al., 2011; Rosenbaum et al., 2011).

ACKNOWLEDGMENTS

We thank the reviewers for providing valuable input. We are particularly grateful to Pawel Andrzej Herman (acting anonymously during the review process) for his numerous constructive comments which helped us to improve the manuscript considerably. This work was supported by the German ministry for education and research (Bundesministerium für Bildung und Forschung, BMBF) within the Bernstein Partner project "Olfactory coding: Integrating experimental and theoretical approaches" (grant no. 01GQ0772 to Randolf Menzel), within the Bernstein Focus "Neural Basis of Learning" (grant 01GQ0941 to Martin Paul Nawrot and Randolf Menzel), and within the Bernstein Center for Computational Neuroscience Berlin (grant no. 01GQ1001D to Michael Schmuker), by the "Innovationsfonds" from Freie Universität Berlin, and by Deutsche Forschungsgemeinschaft within the priority program SPP 1392 "Integrative analysis of olfaction" (SCHM 2474/1-1 to Michael Schmuker).

REFERENCES

Abel, R., Rybak, J., and Menzel, R. (2001). Structure and response patterns of olfactory interneurons in the honeybee, *Apis mellifera*. *J. Comp. Neurol.* 437, 363–383.

Asahina, K., Louis, M., Piccinotti, S., and Vosshall, L. B. (2009). A circuit supporting concentration-invariant odor perception in *Drosophila*. *J. Biol.* 8, 9.

Assisi, C., Stopfer, M., and Bazhenov, M. (2011). Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning. *Neuron* 69, 373–386.

Bhandawat, V., Olsen, S. R., Gouwens, N. W., Schlief, M. L., and Wilson, R. I. (2007). Sensory processing in the Drosophila antennal lobe increases reliability and separability of ensemble odor representations. *Nat. Neurosci.* 10, 1474–1482.

Brandstätter, A. S., and Kleineidam, C. (2011). Distributed representation of social odors indicates parallel processing in the antennal lobe of ants. *J. Neurophysiol.* 106, 2437–2449

Brill, M. F., Reus, I., Rosenbaum, T., Kleineidam, C. J., and Rössler, W. (2011). "Simultaneous recordings from multiple projection neurons in the dual olfactory pathway of the honeybee," *Poster Abstract, in Proceedings of the 9th Göttingen Meeting of the German Neuroscience Society*, Göttingen, T19–31A.

Carlsson, M. A., Chong, K. Y., Daniels, W., Hansson, B. S., and Pearce, T. C. (2007). Component information is preserved in glomerular responses to binary odor mixtures in the moth *Spodoptera littoralis*. *Chem. Senses* 32, 433–443.

Chou, Y.-H., Spletter, M. L., Yaksi, E., Leong, J. C. S., Wilson, R. I., and Luo, L. (2010). Diversity and wiring variability of olfactory local interneurons in the *Drosophila* antennal lobe. *Nat. Neurosci.* 13, 439–449.

Deisig, N., Giurfa, M., Lachnit, H., and Sandoz, J.-C. (2006). Neural representation of olfactory mixtures in the honeybee antennal lobe. *Eur. J. Neurosci.* 24, 1161–1174.

Deisig, N., Giurfa, M., and Sandoz, J. C. (2010). Antennal lobe processing increases separability of odor mixture representations in the honeybee. *J. Neurophysiol.* 103, 2185–2194.

Flanagan, D., and Mercer, A. (1989). An atlas and 3-D reconstruction of the antennal lobes in the worker honey bee, *Apis mellifera* L. (Hymenoptera: Apidae). *Int. J. Insect Morphol. Embryol.* 18, 145–159.

Fonta, C., Sun, X., and Masson, C. (1993). Morphology and spatial distribution of bee antennal lobe interneurones responsive to odours. *Chem. Senses* 18, 101–119.

Galán, R. F., Sachse, S., Galizia, C. G., and Herz, A. V. M. (2004). Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput.* 16, 999–1012.

Galizia, C. G., and Rössler, W. (2010). Parallel olfactory systems in insects: anatomy and function. *Annu. Rev. Entomol.* 55, 399–420.

Galizia, C. G., Sachse, S., Rappert, A., and Menzel, R. (1999). The glomerular code for odor representation is species specific in the honeybee *Apis mellifera*. *Nat. Neurosci.* 2, 473–478.

Ganeshina, O., and Menzel, R. (2001). GABA-immunoreactive neurons in the mushroom bodies of the honeybee: an electron microscopic study. *J. Comp. Neurol.* 437, 335–349.

Grünewald, B. (1999). Morphology of feedback neurons in the mushroom body of the honeybee, *Apis mellifera*. *J. Comp. Neurol.* 404, 114–126.

Hourcade, B., Muenz, T. S., Sandoz, J. C., Rössler, W., and Devaud, J. M. (2010). Long-term memory leads to synaptic reorganization in the mushroom bodies: a memory trace in the insect brain? *J. Neurosci.* 30, 6461–6465.

Huerta, R., Nowotny, T., Garcia-Sanchez, M., Abarbanel, H. D., and Rabinovich, M. (2004). Learning classification in the olfactory system of insects. *Neural Comp.* 16, 1601–1640.

Ito, I., Bazhenov, M., Ong, R. C., Raman, B., and Stopfer, M. (2009). Frequency transitions in odor-evoked neural oscillations. *Neuron* 64, 692–706.

Kazama, H., and Wilson, R. I. (2008). Homeostatic matching and nonlinear amplification at identified central synapses. *Neuron* 58, 401–413.

Kirschner, S., Kleineidam, C. J., Zube, C., Rybak, J., Grünewald, B., and Rössler, W. (2006). Dual olfactory pathway in the honeybee, *Apis mellifera*. *J. Comp. Neurol.* 499, 933–952.

Knudsen, J., Eriksson, R., and Gershenzon, J. (2006). Diversity and distribution of floral scent. *Bot. Rev.* 72, 1–120.

Krofczik, S., Menzel, R., and Nawrot, M. P. (2009). Rapid odor processing in the honeybee antennal lobe network. *Front. Comput. Neurosci.* 2:9. doi:10.3389/neuro.10.009.2008

Laurent, G., MacLeod, K., Stopfer, M., and Wehr, M. (1998). Spatiotemporal structure of olfactory inputs to the mushroom bodies. *Learn. Mem.* 5, 124–132.

Linster, C., Sachse, S., and Galizia, C. G. (2005). Computational modeling suggests that response properties rather than spatial position determine connectivity between olfactory glomeruli. *J. Neurophysiol.* 93, 3410–3417.

Linster, C., and Smith, B. H. (1997). A computational model of the response of honey bee antennal lobe circuitry to odor mixtures: overshadowing, blocking and unblocking can arise from lateral inhibition. *Behav. Brain Res.* 87, 1–14.

Martin, J. P., Beyerlein, A., Dacks, A. M., Reisenman, C. E., Riffell, J. A., Lei, H., and Hildebrand, J. G. (2011). The neurobiology of insect olfaction: sensory processing in a comparative context. *Prog. Neurobiol.* 95, 427–447.

Meyer, A. (2011). *Characterisation of Local Interneurons in the Antennal Lobe of the Honeybee*. Ph.D. thesis, Department of Biology, Universität Konstanz. Konstanz.

Meyer, A., and Galizia, C. G. (2011). Elemental and configural olfactory-coding by antennal lobeneurons of the honey bee (Apis mellifera). *J. Comp. Physiol. A.* doi: 10.1007/s00359-011-0696-8. [Epub ahead of print].

Mobbs, P. G. (1982). The brain of the honeybee *Apis mellifera*. I. The connections and spatial organization of the mushroom bodies. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 298, 309–354.

Moraud, E. M., and Martinez, D. (2010). Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Front. Neurorobot.* 4:1. doi:10.3389/fnbot.2010.00001

Müller, D., Abel, R., Brandt, R., Zöckler, M., and Menzel, R. (2002). Differential parallel processing of olfactory information in the honeybee, *Apis mellifera* L. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 188, 359–370.

Murlis, J. (1997). "Odor plumes and the signal they provide," in *Insect Pheromone Research: New Directions*, eds R. T. Cardé and A. K. Minks (New York, NY: Chapman and Hall), 221–231.

Nowotny, T., Huerta, R., Abarbanel, H. D., and Rabinovich, M. I. (2005). Self-organization in the olfactory system: one shot odor recognition in insects. *Biol. Cybern.* 93, 436–446.

Okada, R., Manz, G., and Menzel, R. (2007). Learning-related plasticity in PE1 and other mushroom body-extrinsic neurons in the honeybee brain. *J. Neurosci.* 27, 11736–11747.

Olsen, S. R., and Wilson, R. I. (2008). Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature* 452, 956–960.

Papadopoulou, M., Cassenaer, S., Nowotny, T., and Laurent, G. (2011). Normalization for sparse encoding of odors by a wide-field interneuron. *Science* 332, 721–725.

Perez-Orive, J., Bazhenov, M., and Laurent, G. (2004). Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *J. Neurosci.* 24, 6037–6047.

Raguso, R. A. (2008). Wake up and smell the roses: the ecology and evolution of floral scent. *Annu. Rev. Ecol. Evol. Syst.* 39, 549–569.

Raman, B., Joseph, J., Tang, J., and Stopfer, M. (2010). Temporally diverse firing patterns in olfactory receptor neurons underlie spatiotemporal neural codes for odors. *J. Neurosci.* 30, 1994–2006.

Riffell, J. A., Abrell, L., and Hildebrand, J. G. (2008). Physical processes and real-time chemical measurement of the insect olfactory environment. *J. Chem. Ecol.* 34, 837–853.

Rosenbaum, T., Brill, M., Rössler, W., and Nawrot, M. P. (2011) "Do antennal lobe output neurons employ a latency code?" *Poster Abstract, in Proceedings of the 9th Göttingen Meeting of the German Neuroscience Society*, Göttingen, T19–30A.

Sachse, S., and Galizia, C. G. (2002). Role of inhibition for temporal and spatial odor representation in olfactory output neurons: a calcium imaging study. *J. Neurophysiol.* 87, 1106–1117.

Sachse, S., and Galizia, C. G. (2003). The coding of odour-intensity in the honeybee antennal lobe: local computation optimizes odour representation. *Eur. J. Neurosci.* 18, 2119–2132.

Sachse, S., Rappert, A., and Galizia, C. G. (1999). The spatial representation of chemical structures in the antennal lobe of honeybees: steps towards the olfactory code. *Eur. J. Neurosci.* 11, 3970–3982.

Schmuker, M., and Schneider, G. (2007). Processing and classification of chemical data inspired by insect olfaction. *Proc. Natl. Acad. Sci. U.S.A.* 104, 20285–20289.

Schmuker, M., Schwarte, F., Brück, A., Proschak, E., Tanrikulu, Y., Givehchi, A., Scheiffele, K., and Schneider, G. (2006). SOMMER: self-organising maps for education and research. *J. Mol. Model.* 13, 225–228.

Seki, Y., Rybak, J., Wicher, D., Sachse, S., and Hansson, B. S. (2010). Physiological and morphological characterization of local interneurons in the *Drosophila* antennal lobe. *J. Neurophysiol.* 104, 1007–1019.

Sigma-Aldrich. (2004). *Flavors and Fragrances*. Milwaukee, WI: Sigma-Aldrich.

Silbering, A. F., and Galizia, C. G. (2007). Processing of odor mixtures in the *Drosophila* antennal lobe reveals both global inhibition and glomerulus-specific interactions. *J. Neurosci.* 27, 11966–11977.

Silbering, A. F., Okada, R., Ito, K., and Galizia, C. G. (2008). Olfactory information processing in the *Drosophila* antennal lobe: anything goes? *J. Neurosci.* 28, 13075–13087.

Stopfer, M., Bhagavan, S., Smith, B. H., and Laurent, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390, 70–74.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.

Strube-Bloss, M. F., Nawrot, M. P., and Menzel, R. (2011). Mushroom body output neurons encode odor-reward associations. *J. Neurosci.* 31, 3129–3140.

Tabor, R., Yaksi, E., Weislogel, J.-M., and Friedrich, R. W. (2004). Processing of odor mixtures in the zebrafish olfactory bulb. *J. Neurosci.* 24, 6611–6620.

Thistle, H. W., Peterson, H., Allwine, G., Lamb, B., Strand, T., Holsten, E. H., and Shea, P. J. (2004). Surrogate pheromone plumes in three forest trunk spaces: composite statistics and case studies. *For. Sci.* 50, 610–625.

Vergassola, M., Villermaux, E., and Shraiman, B. I. (2007). "Infotaxis" as a strategy for searching without gradients. *Nature* 445, 406–409.

Vickers, N. J. (2000). Mechanisms of animal navigation in odor plumes. *Biol. Bull.* 198, 203–212.

Wang, J. W., Wong, A. M., Flores, J., Vosshall, L. B., and Axel, R. (2003). Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* 112, 271–282.

Wick, S. D., Wiechert, M. T., Friedrich, R. W., and Riecke, H. (2010). Pattern orthogonalization via channel decorrelation by adaptive networks. *J. Comp. Neurosci.* 28, 29–45.

Wilson, R. I., and Laurent, G. (2005). Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the *Drosophila* antennal lobe. *J. Neurosci.* 25, 9069–9079.

Wilson, R. I., Turner, G. C., and Laurent, G. (2004). Transformation of olfactory representations in the *Drosophila* antennal lobe. *Science* 303, 366–370.

Wright, G. A., Lutmerding, A., Dudareva, N., and Smith, B. H. (2005). Intensity and the ratios of compounds in the scent of snapdragon flowers affect scent discrimination by honeybees (*Apis mellifera*). *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 191, 105–114.

Wright, G. A., Skinner, B. D., and Smith, B. H. (2002). Ability of honeybee, *Apis mellifera*, to detect and discriminate odors of varieties of canola (*Brassica rapa* and *Brassica napus*) and snapdragon flowers (*Antirrhinum majus*). *J. Chem. Ecol.* 28, 721–740.

Yamagata, N., Schmuker, M., Szyszka, P., Mizunami, M., and Menzel, R. (2009). Differential odor processing in two olfactory pathways in the honeybee. *Front. Syst. Neurosci.* 3:16. doi:10.3389/neuro.06.016.2009

Zollner, G. E., Torr, S. J., Ammann, C., and Meixner, F. X. (2004). Dispersion of carbon dioxide plumes in African woodland: implications for host-finding by tsetse flies. *Physiol. Entomol.* 29, 381–394.

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.

Received: 05 September 2011; paper pending published: 20 September 2011; accepted: 01 December 2011; published online: 28 December 2011.

Citation: Schmuker M, Yamagata N, Nawrot MP and Menzel R (2011) Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee. Front. Neuroeng. ***4****:17. doi: 10.3389/fneng.2011.00017*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 08 February 2012
doi: 10.3389/fneng.2012.00002

Olfactory sensor processing in neural networks: lessons from modeling the fruit fly antennal lobe

J. Henning Proske, Marco Wittmann and C. Giovanni Galizia*

Department of Neurobiology, University of Konstanz, Konstanz, Germany

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Michael Schmuker, Freie Universtiät Berlin, Germany
Kevin Daly, West Virginia University, USA

****Correspondence:***
C. Giovanni Galizia, Department of Neurobiology, University of Konstanz, 78457 Konstanz, Germany.
e-mail: giovanni.galizia@uni-konstanz.de

The insect olfactory system can be a model for artificial olfactory devices. In particular, *Drosophila melanogaster* due to its genetic tractability has yielded much information about the design and function of such systems in biology. In this study we investigate possible network topologies to separate representations of odors in the primary olfactory neuropil, the antennal lobe. In particular we compare networks based on stochastic and homogeneous connection weight distributions to connectivities that are based on the input correlations between the glomeruli in the antennal lobe. We show that moderate homogeneous inhibition implements a soft winner-take-all mechanism when paired with realistic input from a large meta-database of odor responses in receptor cells (DoOR database). The sparseness of representations increases with stronger inhibition. Excitation, on the other hand, pushes the representation of odors closer together thus making them harder to distinguish. We further analyze the relationship between different inhibitory network topologies and the properties of the receptor responses to different odors. We show that realistic input from the DoOR database has a relatively high entropy of activation values over all odors and receptors compared to the theoretical maximum. Furthermore, under conditions in which the information in the input is artificially decreased, networks with heterogeneous topologies based on the similarity of glomerular response profiles perform best. These results indicate that in order to arrive at the most beneficial representation for odor discrimination it is important to finely tune the strength of inhibition in combination with taking into account the properties of the available sensors.

Keywords: olfaction, model, antennal lobe, inhibition, odor separation

1. INTRODUCTION

The design of artificial olfactory systems can benefit greatly from a comparison with natural olfactory systems in animals. For example, the olfactory system of insects has an evolutionary history of more than 400 million years, resulting in a highly efficient and functional network of neurons that process olfactory information. To make use of this evolutionary design knowledge, we need to understand the interaction between the characteristics of olfactory receptors and the processing in the first olfactory neuropil, the antennal lobe (AL).

Drosophila melanogaster is an ideal model system to study this interaction. The available genetic tools have allowed the detailed and extensive study of both olfactory receptor neurons and the architecture of downstream processing networks. The processing stream in *Drosophila* starts with a number of olfactory receptor neurons (ORNs) on the antenna that transduce chemical stimuli into neural signals. Each ORN expresses a few olfactory receptor types and then projects to one glomerulus in the AL where ORNs synapse onto projection neurons (PNs; Vosshall et al., 1999; Gao et al., 2000; Bargmann, 2006). Receptor neurons of the same type expressing the same receptor mostly project to only one glomerulus (Vosshall et al., 1999). The PNs then carry information from the glomeruli to higher brain areas like the mushroom bodies or the lateral horn (Stocker, 1994). Glomeruli can be identified on anatomical and physiological optical sections of the AL and it has therefore been possible to establish the spatiotemporal response patterns specific to particular odor stimuli (Fiala et al., 2002; Wang et al., 2003; Silbering and Galizia, 2007; Silbering et al., 2008). The activation patterns in the AL are not only caused by stimulus-dependent receptor activation but also by secondary lateral activation within the AL. This internal AL processing is mediated by local interneurons (LNs) that are activated by ORNs and PNs and influence the activity of other ORNs and PNs (Wilson and Laurent, 2005). *Drosophila* harbors different subsets of LNs both inhibitory as well as excitatory (Olsen et al., 2007; Shang et al., 2007), with the inhibitory type being more prominent (Wilson et al., 2004; Kazama and Wilson, 2008).

The LNs in the AL of *Drosophila* and other insects play an important role in the processing of odor information. For example in *honeybees* it has been shown that the specific heterogeneous connectivity patterns of LNs in the AL are able to decorrelate similar odor representations and hence facilitate odor identification (Sachse and Galizia, 2002). This type of connectivity has been analyzed in detail (Linster et al., 2005). In *Drosophila* homogenous global inhibition is beneficial for the processing of odor stimuli by means of stimulus normalization that allows for efficient encoding of odorant stimuli at different concentrations (Kazama and Wilson, 2008; Asahina et al., 2009) and equalization that maximizes

the mutual information of stimuli and representation in downstream processing areas (Luo et al., 2010; Olsen et al., 2010; Satoh et al., 2010). These results also fit well with evidence from anatomical studies in which in *Drosophila* inhibitory LNs have dense connections throughout the antennal lobe (Stocker et al., 1990; Olsen and Wilson, 2008).

In this study we investigate the interaction between the distribution of receptor response profiles and the processing in different AL network architectures. In particular we focus on the ability of a processing network to make odor representations easier to discriminate. If, in a simplified model, one assumes linear interactions between glomeruli it is possible to represent AL networks as matrices and the resulting processing as a dot product of this matrix with the input vector. The input to our model network comes from the DoOR database (Galizia et al., 2010) which interpolates real physiological response data into relative activation values over a large bank of test odors. This allows us to integrate the odor-response data from the *Drosophila* olfactory system into our model.

We start off by analyzing the odor separation performance in AL networks derived from basic connectivity principles. In a simplified example we show the relationship between excitation and inhibition for these networks. Performance is evaluated by looking at the pair-wise angular separation between odor vectors. This generic description of separation of representations makes reasonable assumptions about the downstream readout networks. We show that with the default response profile distributions of the receptors in the DoOR matrix best performance is achieved with global inhibition between glomeruli and moderate connection weights. Lastly we systematically change the entropy of the input matrix while simultaneously evaluating the performance of our model networks. Under conditions with relatively low entropy in the input matrix connectivity that is based on input correlations is more successful in separating odors than other networks.

2. MATERIALS AND METHODS

2.1. LINEAR TRANSFORMATIONS IN THE AL

Drosophila with its roughly 50 glomeruli (Laissue et al., 1999) and the ease of targeted genetic manipulations is an ideal animal for detailed cellular and molecular analyses. It thus offers a lot of detailed information that can be used for computational models. However, even in this very limited system the number of free parameters can quickly become unfeasible to chart. Also, neurons in the AL have a complex structure and are most likely not straightforward to model with Hodgkin-Huxley equations derived from the squid giant axon. In this study we therefore made the following simplifying assumptions:

1. We neglected the temporal structure of odor stimuli and represented each odor as a static N-dimensional vector of activation values where N corresponds to the number of measured receptor neurons.
2. We assumed that transfer of activity between ORNs and PNs in the glomeruli has a linear input/output relationship and that interactions with other glomeruli by means of excitatory and inhibitory local neurons are also linear.
3. We set all diagonal values in the connectivity matrix to 1. This implies that every olfactory receptor neuron activates its directly efferent projection neurons with the same strength. If there are no intra-AL connections the input vectors get multiplied with a unity matrix and hence the input equals the output.
4. We assume that an activation of zero represents a completely silent cell. Hence we set all negative activation values to 0.

Under these assumptions it is possible to represent the transformation of olfactory neuron activation to PN output as a single connectivity matrix that incorporates all intermediate steps of processing. Let $\mathbf{x}$ be the input vector of olfactory receptor activations and $\mathbf{y}$ be the output of PN activations, where $\mathbf{h}$, $\mathbf{g}$ are activations of intermediate interneuron populations (only two of them considered for simplicity, without loss of generality). These neuron populations have size n_x, n_y, n_h, and n_g, respectively. The connectivity between population α to population β can be expressed in terms of a matrix $W^{\alpha\beta}$ of size n_α-by-n_β. In this matrix the element $W_{i,j}$ maps the activation of unit i in the source population to unit j in the target population (**Figure 1**).

Then the output of this feedforward network can be characterized as:

$$\mathbf{y} = \mathbf{x}W^{xy} + \mathbf{h}W^{hy} + \mathbf{g}W^{gy} \tag{1}$$

where $W_{\alpha\beta}$ is the connectivity matrix from population α to population β of size n_α-by-n_β. The non-$\mathbf{x}$ terms can be described within a linear regime as

$$\mathbf{g} = \mathbf{x}W^{xg} \tag{2}$$

$$\mathbf{h} = \mathbf{x}W^{xh} + \mathbf{g}W^{ih} \tag{3}$$

$$= \mathbf{x}W^{xh} + \mathbf{x}W^{xg}W^{gh} \tag{4}$$

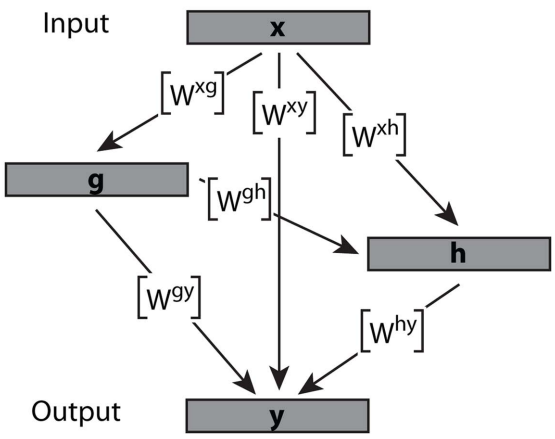

FIGURE 1 | Example of a hierarchical AL network in which the interactions between neuron populations can be reduced to a single weight matrix. The activation of neuron subpopulations is contained within vectors **x**, **g**, **h**, and **y**. Ultimately the combined matrices map the input **x** onto the output vector **y**.

Substituting the terms from equations 2 and 3 into equation 1 results in

$$\mathbf{y} = \mathbf{x}W^{xy} + \mathbf{x}W^{xh}W^{hy} + \ldots \\ \mathbf{x}W^{xg}W^{gh}W^{hy} + \mathbf{x}W^{xg}W^{gy} \tag{5}$$

By means of distributivity

$$\mathbf{y} = \mathbf{x}\left(W^{xy} + W^{xh}W^{hy} + \ldots \\ W^{xg}W^{gh}W^{hy} + W^{xg}W^{gy}\right) \tag{6}$$

where all matrix multiplications inside the brackets result in matrices of size n_α-by-n_β and can hence be added to give one single connectivity matrix. This simplification holds for all linear feedforward networks.

The anatomy of the first-order olfactory relay in insects can hence be approximated by a single compound connectivity matrix that represents the functional interactions between glomeruli (**Figure 2A**). This matrix provides a linear map from the input x to PN output y:

$$y_j = \phi\left(\sum_i w_{ij}x_i\right) \tag{7}$$

where y_j is the activity of a projection neuron, x_i is the activity of an olfactory receptor neuron in response to an odorant stimulus, w_{ij} is the connection weight from ORN i to PN j, and ϕ is the activity transfer function. This is equivalent to the dot product between a vector and a matrix in linear algebraic terms. Conveniently the vector elements correspond to the activation of ORNs and the weight matrix represents the functional interactions between glomeruli (**Figure 2B**). In our simulations the dimensionality of x and of y are identical, reflecting the natural situation in insects, where the number of glomeruli corresponds to the number of receptor neuron input families, an where uniglomerular projection neurons of a single glomerulus share their response properties (Hallem and Carlson, 2006).

The activation values from the DoOR input matrix are scaled between 0 and 1 where 0 corresponds to the neuron being silent. Similarly we implement a lower limit for the projection neurons

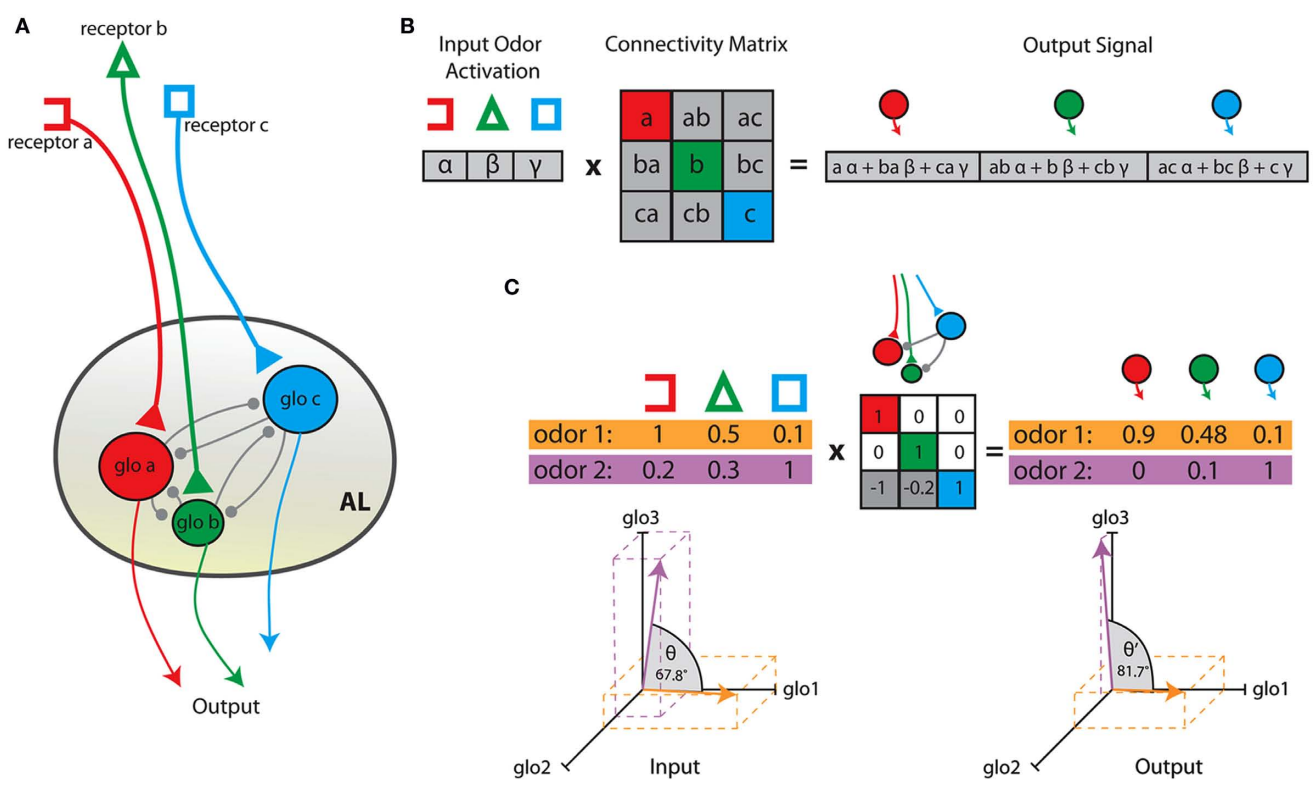

FIGURE 2 | Schematic of the relationship between antennal lobe anatomy and the linear transformation model. (A) Shows a simplified antennal lobe architecture. Different receptor types innervate one glomerulus exclusively (color-coded). Local interactions (gray) between glomeruli are mediated by local interneurons. The output from the AL is the activation of PNs in each glomerulus. **(B)** Odors are represented as vectors of receptor activation and can be transformed into PN output by means of multiplication of the input vector with a connectivity matrix. Diagonal values correspond to a gain factor on the activation of each glomerulus whereas non-zero off-diagonal values lead to interglomerular interactions. **(C)** Shows an example multiplication for two odors and a hypothetical AL network with inhibitory connections from glomerulus c to a and from glomerulus c to b. Negative output neuron activity is set to 0. The diagrams show an increase in angular distance in 3D space between the two odor vectors before (left) and after (right) AL processing.

creating linear threshold dynamics:

$$\phi(x) = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad (8)$$

2.2. INPUT DATA

As input data we used the DoOR database of odorant responses (Galizia et al., 2010). This database uses a sophisticated interpolation algorithm to combine odorant responses from different studies and provides response values from 50 receptor types to 250 odorant stimuli (**Figure 3A**). In order to decrease the amount of empty values in the input matrix we excluded all receptors that (a) did not have a known glomerular target, and (b) had recorded responses to less than 70 odorants. We also excluded all odorants with responses measured in less than 8 receptor types. This left us with an odorant receptor response matrix with 137 odorant responses in 22 receptor types. All missing values (ca. 10%) were filled in with the spontaneous activity values. All receptor types in our input matrix had a known one-to-one correspondence with one AL glomerulus so that we could assume the receptor activation to be equal to the input activation in the corresponding glomerulus.

2.3. GENERATION OF CONNECTIVITY MATRICES

We generated connectivity matrices according to certain design principles previously hypothesized to play a role in the interglomerular interactions in *Drosophila* and other insects (e.g., Linster et al., 2005). In particular, we were interested in the question if networks based on external factors such as response similarity or spatial position in the antennal lobe provide advantages over random networks of similar connection strength distributions. To this end we calculated a matrix of correlation coefficients of size 22 by 22 from the input matrix. There are several possible ways to translate correlation coefficients into connection weights. The main question is how to interpret negative values (indicating anti-correlation between input vectors of two glomeruli). We set all negative coefficients in the matrix to zero and then normalized all values to fall in the range from 0 to 1. We also tested alternative scenarios, e.g., to scale all correlation values or to translate anti-correlation into excitatory connections, with only minor quantitative effects (data not shown). As a last step the diagonal values were all set to 1 in order to simulate the direct input from the ORNs to the PNs in the glomeruli. This resulted in a 22×22 connection matrix with unity value on the diagonal and otherwise negative values. Glomeruli with highly correlated input profiles were thus set to have inhibitory influence on each other, proportional to their input correlation coefficients (**Figure 3Bi**). In order to test to what extent this correlation-based matrix is superior to a random matrix we scrambled the values in the original matrix, keeping the diagonal and overall symmetry intact (**Figure 3Ci**). Repeating this step with different random seeds generated 50 control matrices

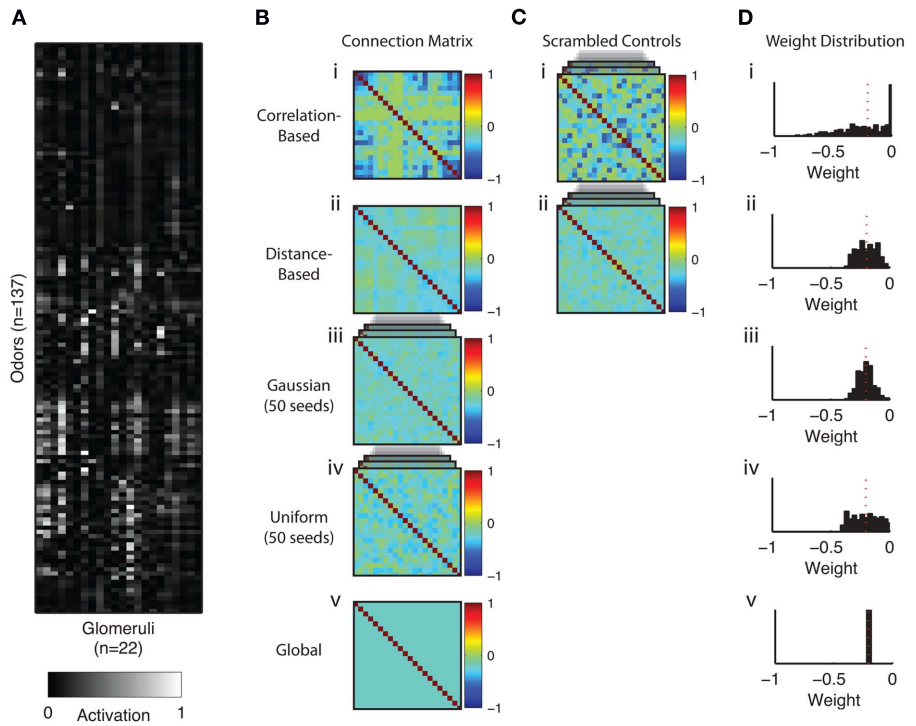

FIGURE 3 | Input data and connection matrices derived from different schemes. (A) Activation values in the DoOR input matrix. **(Bi–v)** AL connectivity matrices. **(C)** Scrambled control matrices (50 different random seeds). **(D)** Distribution of connection weights in the matrices. All matrices were normalized to have the same mean of the off-diagonal values as the correlation-based matrix. This is indicated by the red dotted line. The positive diagonal values are not shown in the histograms.

with the same distribution of connection weights as the original (**Figure 3Di**).

In a second step we also generated matrices based on the spatial location of glomeruli within the antennal lobe. We calculated the Euclidean distance between the center of the glomeruli from a 3D reconstruction of a whole AL (Laissue et al., 1999), normalized these values to fit on a range between 0 and 1. Again, the diagonal was set to 1 (**Figure 3Bii**). We generated 50 control matrices as described above (**Figure 3Bii**). There are in fact many ways in which the spatial distances between glomeruli can be assessed, e.g., taking into account the connectivity paths of neurites in the antennal lobe. Here we restricted ourselves to a simple Euclidean distance assuming that the resulting connectivities form a generic, if not necessarily accurate, representation of distances in the AL. It should be noted that in the fruit fly, glomeruli fill the antennal lobe in the entire volume. The situation is different to that found in mammals or in honeybees, where glomeruli form a hemispheric sheet. In the latter situation, geometric distance can be calculated in multiple ways (along the sheet, through the center of the AL, etc.). This does not apply for the fruit fly.

Two unrelated types of matrices were constructed by sampling connections from a Gaussian and uniform probability distribution. Fifty matrices were generated with different random seeds (**Figures 3Biii,iv**) and the diagonal set to 1.

Lastly we tested the performance of a matrix approximating an architecture of global inhibition mediated by a single inhibitory neuron. For this connectivity matrix we fixed all off-diagonal weights to one value equivalent to the mean off-diagonal weight of the correlation-based matrix (**Figure 3Bv**). The diagonal was again set to 1.

All resulting networks were normalized to have the same mean weight value (see **Figure 3C**). The overall efficacy of the inhibitory connections was changed by multiplying all off-diagonal values with a scaling factor. This factor was the only free parameter in our model. A negative scaling factor is hence equivalent to an inhibitory network and a positive scaling factor to an excitatory network.

2.4. QUANTIFICATION OF SEPARABILITY

To benchmark the performance of different AL networks we measured the geometric angle between odorant response vectors. We assumed that two odorants are more easily separable if the angle between their response vectors is larger. The performance P of a network was quantified as the mean of the sine of the angles between all odors.

$$P = \frac{2}{n^2 - n} \sum_{i=0, i<j}^{n} \sin \alpha_{ij} \tag{9}$$

where P is the performance index and α_{ij} is the angle between odors i and j. The term before the sum normalizes the result by the number of angles between the vectors, taking into account only the top half of the angular distance matrix. Taking the sine of the angles has the advantage of promoting a large number of smaller angles over a few large angles in the performance index. **Figure 2C** shows an example of how input data and connectivity matrix interact to change the angle between two arbitrary vectors in two-dimensional space.

In addition, we sought to evaluate the metabolic efficiency of the tested networks. Since the possible activity value in each glomerular unit was limited to positive values strong inhibition that resulted in negative activation was wasted from a metabolic stand point. We hence evaluated the efficiency E of representation by first calculating z_{jk} as the PN activations to odor k without setting negative values to 0:

$$z_{jk} = \sum_{i,k} w_{ij} x_{ik} \tag{10}$$

and then calculating E as the mean negative overshoot over all odors for each network:

$$E = \frac{1}{\sum_{j,k} \left[z_{j,k} < 0 \right]} \sum_{j,k} z_{jk} \left[z_{jk} < 0 \right] \tag{11}$$

using Iverson brackets for the conditional sum. E is hence 0 in the most efficient case where there is no overshoot. The more negative the value for E the less efficient the network. Similarly, the sparseness S of a representation was estimated by calculating the proportion of silent glomeruli over all odors.

$$S = \frac{1}{n_{jk}} \sum_{j,k} \left[y_{jk} = 0 \right] \tag{12}$$

where $y_{j,k}$ are the PN activations as in equation 7 but for all odors, and n_{jk} is the number of elements in the output matrix (22 glomeruli × 137 odors).

Input data from the DoOR database was prepared in R. All simulations were carried out using MATLAB (Natick, MA, USA).

2.5. CHANGING INPUT ENTROPY

Lastly we were interested in how the specific distribution of receptor activations present in the DoOR matrix interacts with the network models described above. We hypothesized that networks based on relatively homogeneous inhibition would perform well in scenarios in which response profiles for different odors are well-separated in the input by their strongest responsive glomeruli. If, however, the input is structured so it is often the same glomeruli that is the most active across odors, homogeneous inhibition will not be able to implement a winner-take-all mechanism as efficiently. In order to evaluate this systematically we calculated a measure for the input matrix that is related to the Shannon entropy (Shannon, 1948) of the responses of each glomerulus over all odorants in the database, H_{glo}:

$$H_{glo} = -\sum_{i=1}^{n} p(x) log p(x) \tag{13}$$

where n is the number of glomeruli used, 22, and $p(x)$ is the probability of seeing a particular activity value in this glomerulus. Instead of binning the response values we transformed them into a rank order code by first assigning a number to each odor

response in the input matrix so that for each odor stimulus the most active glomerulus was 1, the second most active glomerulus was 2 and so on. The total entropy of an input matrix was then calculated as

$$H_{tot} = \sum_{j=1}^{m} H_{glo} \tag{14}$$

where m is the number of odors in the matrix. The original DoOR input matrix had an H_{tot} of 57.6. The theoretical maximum of our measure value distribution of a 137 × 22 input matrix would be 68.0 if activity values were optimally distributed over glomeruli for all odors. Interestingly the original DoOR derived input matrix scores quite high on our entropy score, showing that over the tested odors glomerular responses are distributed quite homogeneously.

We generated input matrices with different H_{tot} by taking the original DoOR matrix and progressively ordering individual odor vectors to have the first glomerulus responding the strongest, the second glomerulus the second strongest, and so on. The extreme case where all odor vectors are ordered for response strength then has an H_{tot} of 0.

3. RESULTS

3.1. SIMULATIONS IN A TWO-DIMENSIONAL SYSTEM

We first asked what network types are more successful in separating odors in our simple vector model. We started by looking at the angular separation of two arbitrary vectors (**Figure 4A**) in two-dimensional space. This is equivalent to 2 odors being represented by 2 hypothetical glomeruli and the interaction between the glomeruli is the algebraic equivalent of a multiplication of the vectors with a shear matrix. An example with 2 different sets of parameters for *m* and *n* shows an increase of angular separation with negative and a decrease of separation with positive values (**Figure 4B**). **Figure 4C** shows the increase in angular separation between the two vectors shown in **Figure 4A** as the off-diagonal values *m* and *n* are altered systematically from negative to positive values. The highest increase in angular distance between the two vectors is in the lower left corner approaching point 2 with negative values of *m* and *n*. This indicates that inhibition will be particularly efficient in separating the angular distance between odor vectors in contrast to lateral excitation. If $m, n \in \mathbb{R}^+$ as in point 1 the angle between the two vectors decreases toward 0.

In our networks we limit the activation values of our model neurons to values between 0 and 1. If these boundary conditions are implemented with the two example vectors a similar pattern emerges as before (**Figure 4D**). It becomes apparent that the dependence of angular separation on the parameters *m*/*n* in this "cutoff" condition (**Figure 4E**) is similar to the general case (**Figure 4C**) in the sense that the best separation is achieved with negative values *m* and *n*.

The first analysis of the two-vector scenario does not easily allow us to postulate that inhibition is necessarily the best way to achieve angle separation in the realistic case with 22 dimensions and 137 odor vectors. The following analysis is therefore based on numerical simulations with real input data from the DoOR database.

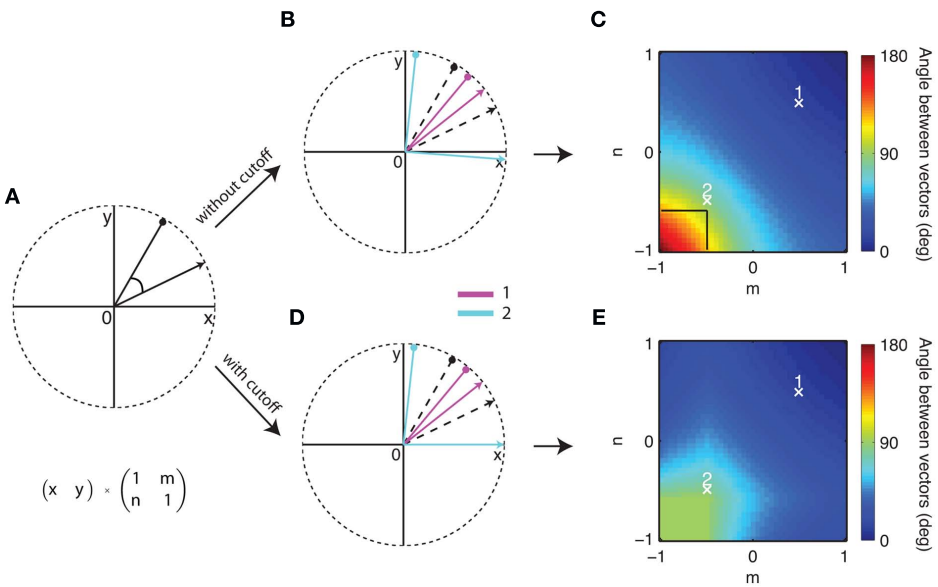

FIGURE 4 | The effect of a two-dimensional shear matrix on two arbitrary vectors (A). (B) Shows an example of the separation of the input vectors in **(A)** – black dashed lines – at two points in *m, n* space. Negative values of *m, n* separate the vectors (cyan) whereas positive decrease the angle between them (magenta). **(C)** Shows the angle between the example vectors from **(A)** over a range of *m* and *n* from −1 to 1. The example values from **(B)** are marked with white crosses. In general, the more negative *m* and *n* the larger the angle of separation between the two vectors becomes. In the **(D,E)** all negative values in the output vector are set to zero restricting the maximum separation that can be achieved to 90°. For a direct comparison the black box in **(C)** indicates the maximum separation in **(E)**.

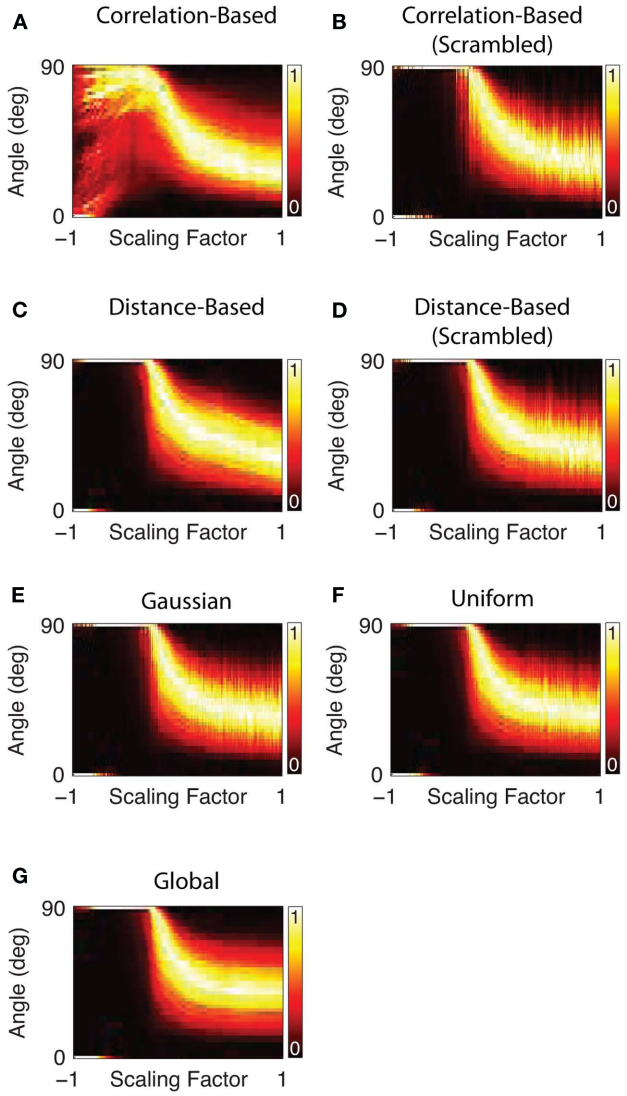

FIGURE 5 | Angle separation between all odor pairs in different AL network models. Each figure shows the distribution of pairwise angles between odor vectors (false-color coded) as a function of the scaling factor. For each scaling value the resultant distribution of angles is normalized so that the most probable angle is 1. A positive scaling factor indicates excitation and a negative scaling factor inhibition. The flip from the condition where all angles are 90–0° is an artifact of the very small vector lengths due to high negative scaling (see **Figure 6**).

3.2. SIMULATIONS WITH REAL-WORLD INPUT DATA

We next asked how different network architectures performed in separating the input odor vectors for a potential downstream readout mechanism in a multidimensional system. We measured the angular distance between all odor pairs in these different conditions. We assumed that in general a larger Euclidean angle between two odor vectors will make the two underlying odors more easily separable. Note, that we did not make any explicit assumptions about the exact form of a readout mechanism.

This multidimensional analysis confirmed that in general lateral inhibition pushes the distribution of angles between odor vectors up toward the maximum of 90° (**Figure 5**). If the scaling parameter is positive (i.e., the lateral interaction is excitatory) the angles between odors decrease, pushing individual vectors closer together, and making them more difficult to separate by a readout mechanism. It can be observed that with moderate inhibition the natural distribution of angles between odor vectors is shifted toward larger angles. This is due to the resulting sparsification of the odor vectors. Small glomerular activations get pushed to zero by strong lateral inhibition. Geometrically, this means that an increasing number of angles between odors become perpendicular in glomerulus space. At a scaling factor around the value of −0.5 many vectors have become perpendicular and the distribution is not visible any more. However, this comes at the cost that some odor representations end up with an identical representation in the output. With very high negative scaling an increasing number of vectors effectively have zero length and cannot be distinguished from each other (**Figure 6A**). The number of vectors that are effectively identical starts to increase with scaling values lower than −0.5 (**Figure 6B**). This image is the same for all network types tested with the notable exception of the correlation-based network. This network type has a much less homogeneous development of the distribution of angles between odor vectors with high negative scaling factors (**Figure 5A**). This is mirrored in the minimum vector length which is generally higher than in the other networks for scaling factors below −0.5 (**Figure 6A**). With such high negative scaling the correlation-based network also has an advantage when comparing the number of identical vectors after processing to the other network types (**Figure 6B**).

A systematic comparison of the separation performance of different types of network using the performance index in equation 9 is shown in **Figure 7**. Again moderate inhibition was able to separate odor representations whereas excitation and strong inhibition decreased the separation performance. In our model with the natural DoOR matrix input, networks derived from input correlations performed worse in separating odor vectors from each other than random networks, globally inhibited networks, or the network derived from the spatial relationships between glomeruli in the AL. Note, that this is not due to the heterogeneous distribution of connection weights (see **Figure 3Di**) as the scrambled version of the correlation-based matrix performs similarly to the other networks.

Another important factor is the metabolic efficiency with which a system is able to perform its tasks. We therefore asked to what extent each of the networks produces inhibition that exceeds the necessary amount for silencing the glomeruli. The efficiency as calculated in equation 11 mirrors the results from network performance. It is near identical for all network types apart from the correlation-based network which is slightly less efficient in this measure.

To explain why the performance and angle distributions of the correlation-based network are different from the other networks we next looked at how sparse the odor vectors become after AL processing. We hypothesized that in the input-correlation-based network strong glomeruli tend to inhibit each other selectively and prevent the winner-take-all mechanism inherent in the more

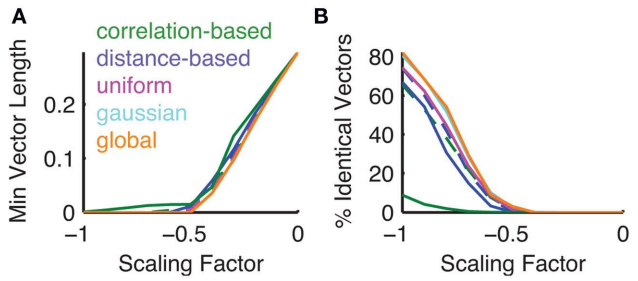

FIGURE 6 | Analysis of vectors after processing by different network types. (A) The minimum vector length for all network types apart from the correlation-based network shows a decrease to 0 at scaling values around −0.5. Dashed lines indicate mean results over 50 seeds for the scrambled versions of the correlation-based (green) and distance-based networks (blue). **(B)** Similarly the number of identical vectors that cannot be distinguished increases drastically at scaling values more negative than 0.5. This effect is attenuated in the correlation-based network.

homogeneously inhibited networks. If this was the case then we would expect that sparseness as defined by the number of silent glomeruli (equation 12) would be lower than in the other networks. **Figure 7C** shows the relationship between the sparseness S and the scaling factor. Indeed, the correlation-based network produces noticeably less sparse representations of odors than the other networks. This becomes even more apparent when looking at the representation of an example odor after processing. **Figure 7D** shows the relative activation values before (input) and after processing by two networks of the correlation-based type and its scrambled control. While both the correlation-based and scrambled correlation-based network generate sparse representations this sparseness is more heterogeneous in the correlation-based network, i.e., the amount of input activation is no reliable predictor for the resulting output activation in this network. Due to the stochasticity of the scrambled network the weakest inputs also lead to the weakest outputs. Because the amount of effective inhibition is the same in both networks the heterogeneous output in the functional network case interacts with the zero-cutoff threshold to result in less sparse networks. This is most visible in the distribution of angles after processing (**Figure 5A**).

3.3. DIFFERENT INPUT ENTROPIES

In order to compare the performance of different AL network types under changing input statistics we first generated a number of networks with different Shannon entropies. We determined the best possible network instance for each network type by sweeping over a range of negative scaling factors and then plotting only the best possible performance against the entropy of the input network (**Figure 8**). The random network types (scrambled controls, Gaussian, and uniform) were tested with 10 different random seeds and did not show large SEM. We also tested two different types of correlation-based networks: one that was based on the original input data (correlation-based) and one where the correlation-based weights were recomputed for each instance of the input data (correlation-based flexible).

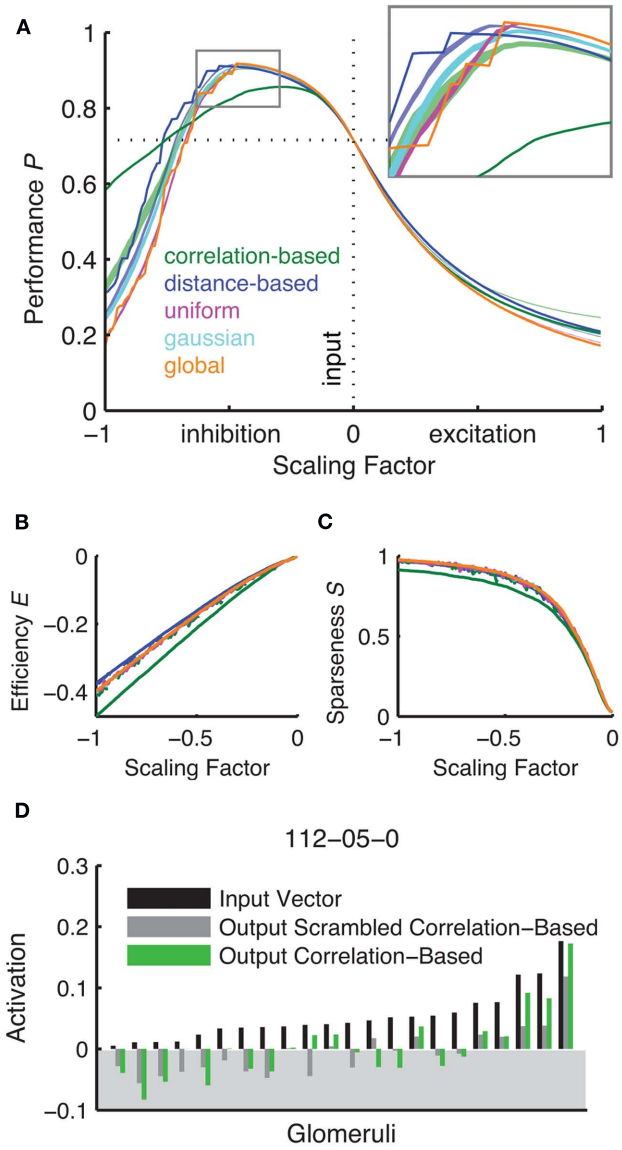

FIGURE 7 | Inhibitory networks lead to an increase in angular separation between odor pairs. A peak in performance can be observed with moderate inhibition **(A)** as compared to strong inhibition and excitation. The inset shows a magnification of the different network types' performance around the peak. Unity connection matrix performance is indicated with the dotted horizontal line. The random networks are represented by shaded areas of 50 seeds ±1 SEM (green for the correlation-based scrambled network and blue for the distance-based scrambled network). **(B,C)** Show the efficiency and sparseness of the resulting transformation against different scaling factor values. Here the mean of 50 different scrambled version is shown as a dashed line in green for the correlation-based and blue for the distance-based network. Positive scaling factors are excluded as they do not produce overshoot/sparseness. **(D)** Shows example odor representations before and after AL processing for nonaic acid at a scaling factor of 0.25, negative activity would be set to 0 (shaded area).

At first, the performance in all networks increases slightly (**Figure 8**). The input data set contains some correlations and

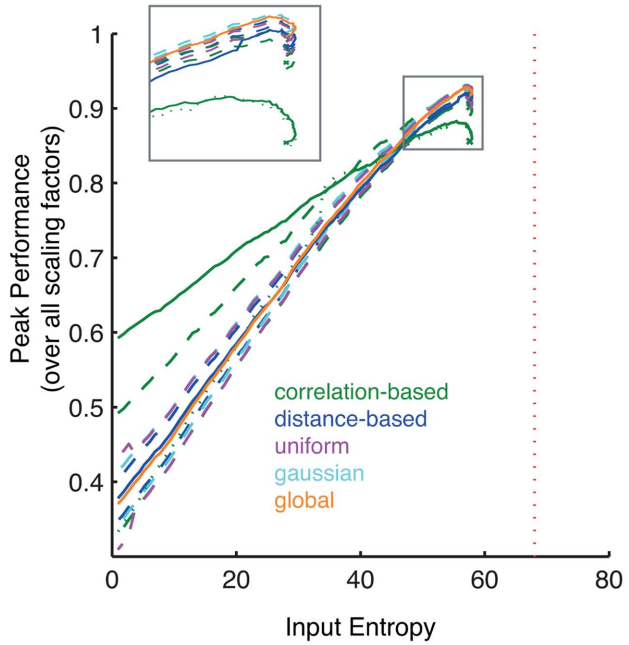

FIGURE 8 | Change of peak performance with different constructed input entropies. The lower the entropy the less random the input odor matrix. The Peak Performance refers to the mean sine angle differences as in **Figure 7A**, tested over a range of negative scaling factors. The vertical red dotted line is the theoretical maximum. The crosses represent the peak performances of the input from the original DoOR matrix. Dashed lines indicate the mean $\pm$ 1 SD for the random networks, testing 10 different random seeds. The dotted line is a correlation-based network that has been recalculated with each new instance of the input data. With increasing entropy the peak performances decrease. The inset shows a magnification of the peak performances with high entropies. Note, that the lower SD line for the scrambled correlation-based networks is mostly hidden underneath the line for the uniform networks.

therefore has entropy below the theoretical maximum. These inherent correlations are reduced when odor vectors are ordered, leading to the initial increase in entropy.

After the transient increase the peak performances decrease drastically: input with less information leads to less information in the output. Note, that the peak performance values drop below their initial values at lower entropies levels than in the original data. This asymmetry is based on the entropy formula (equation 13). At first the entropy increases because individual rank order values in a given response profile become more likely. As more and more values become identical though, the log term increases non-linearly.

The peak odor-response separation performance of each network type is thus highly dependent on the total entropy of the input matrix (as measured by our index). While the correlation-based network performs worse than the other networks at the entropy of the DoOR input it performs better than other networks with input matrices of lower entropy. This can not be attributed to the heterogeneous distribution of connection weights in this matrix, as its scrambled control also falls below the performance of the unscrambled network. Surprisingly, if the correlation-based network is recomputed with each new instance of the input this network type loses its advantage at low entropies.

A different pattern emerges with the global inhibition network with decreasing entropy. Under high entropy conditions this network type performs generally better (**Figure 8**, inset) but performance decreases more drastically under low input entropy conditions. The networks based on random distributions have performances in between these two network type. These results emphasize the importance of the interaction between the statistical properties of the input and the AL network processing. Studying both independently does not reveal the full picture.

4. DISCUSSION

Here we have presented a model of the *Drosophila* antennal lobe input-output transformation under the assumption that the topological effect of the network at this level in the odorant processing stream can be approximated by linear functions. This facilitates both the modeling and analysis and provides an intuitive understanding of underlying interactions. Thus while our model – as every simplifying model – can analyze some critical features of the AL network, it can not be taken as an explanation of the natural situation. It can, however, be used as a guide for artificial olfaction, where simplified algorithms add to the efficiency.

In particular, our results help to understand the role that inhibitory interactions play in influencing separability of odorant stimuli as measured by the average pair-wise angular distance between odor representations. To start with, the simplification of the problem into a two-dimensional space allows us to appreciate the connectivity between glomeruli and the resulting interactions as the off-diagonal values in a shear-type transformation matrix. We showed that both in an unbounded scenario and in a scenario where only positive activations are allowed negative shear values separate vectors inside the space (**Figure 4**). This underlines the robustness of this modeling approach. Positive shear values push vectors together with the extreme scenario of making the vectors all point in the same direction.

Such a simplified description can only serve to highlight the key principles of linear interactions in the AL. It is important to also take into account both the statistics of real input data. One aspect not analyzed here is the interaction of v vectors in n-dimensional space where $v < n$, i.e., when the number of odors that need to be coded is much higher than the number of available odor sensors (glomeruli). We anticipate that in this case the optimal solution is likely to be based on the exact distribution of the input vectors in space, as suggested by our entropy analysis (**Figure 8**, and see below). A study by Schmuker and Schneider (2007) used a large input dataset generated from the chemical properties of more than 800 odorants and found similar winner-take-all dynamics in a correlation-based lateral inhibition network under the assumption that receptive fields in olfactory sensory neurons are indeed equivalent to the ones generated in the study.

The high level of abstraction in our model has some advantages over more classical approaches. The simplicity of the underlying equations leads to a relatively small number of free parameters.

A neural network model with close-to-realistic neurons needs to model the electrical properties of the neurons (Bazhenov et al., 2001; Linster et al., 2005). A single Hodgkin-Huxley neuron model

with only sodium, potassium, and leak currents will already require at least 19 parameters to be tuned (maximum conductance and reversal potential plus 4 parameters for activation and inactivation variables for each current, plus membrane capacity; Hodgkin and Huxley, 1952). As a consequence, a simplification of the theoretical description of such networks has already been implemented in other studies (Luo et al., 2010; Papadopoulou et al., 2011). We believe that our simplified model with linear threshold neurons and only one free parameter (scaling factor) and one set of assumptions (the connectivity values in the matrix) is the most parsimonious model that allows to draw conclusions about the network topology. In addition, the resulting speed of simulation allows for the efficient coverage of large parts of the parameter space. In this way it was possible to map the relationship between glomerular interaction and network performance exhaustively.

For the networks that we tested we observed a common pattern. As predicted by the reduced two-dimensional model, inhibition led to vector separation whereas excitation led to a reduction of angles between odor vectors. This has potential advantages for the reliability of recognition of certain odors but comes at the cost of a lower capacity for odor representation when the number of available coding dimensions is much lower than the number of odor identities the fly needs to be able to distinguish. It is not known what this number is in flies (i.e., how many odors need to be distinguished ecologically); in our simulation the number of odors was about 6-fold (137 odors for 22 glomeruli). It should be noted that the antennal lobe does not only compute odor discriminability. Odor identity (as an interplay of generalization and discriminability) and concentration invariance are but two of several additional tasks likely to be relevant in the antennal lobe network. It is feasible that excitatory interneurons play an important role here. The processes observed here potentially allow for a graded mechanism in which modulatory connections attenuate interglomerular inhibition and hence the sparseness of AL output. This could serve to adjust the reliability-capacity trade-off to the requirements of the animal's environment or life stage. We also found that in all networks inhibition strength has a sweet spot, suggesting that biological systems carefully adjust the amount of inhibition that is exerted in their networks. With increasing inhibition the network degenerates, i.e., several odors collapse onto the same output vector. In addition, output vectors become too short, thus losing robustness against noise.

All networks were subject to a soft winner-take-all mechanism. The network based on input correlations exhibited a more complex behavior when subjected to strong interglomerular inhibition and the winner-take-all effect was attenuated. This is due to the fact that the inhibition in the correlation-derived connectivity matrices focuses on the glomeruli that have the highest correlations. This biases the inhibition toward active glomeruli and can leave less active glomeruli that have uncorrelated response profiles unaffected. The separation of odor vectors by soft-winner-take-all dynamics can thus only occur in part in this type of network. This effect is no consequence of the heterogeneity of the connection weights in the correlation-based networks, as shown by the fact that the scrambled version of the network (with the same distribution of weights) performs like the stochastic versions. One important side-effect is that the correlation-based network is less sensitive to non-optimal scaling factors at the cost of lower peak performance – that is performance degrades more gracefully in correlation-based networks.

We compared odor separability by measuring the angle between all available odor vectors. A maximum odor separation capacity would result from a maximization of the angles between all the odors in the high-dimensional coding space. This measure does not make precise assumptions about the possible readout mechanism nor does it include a measure of ecological relevance for the animal. While in our benchmark all odors are equally important it is likely that the biological system places emphasis on odors that are frequently encountered or ecologically very important. Further, with some odor families, generalization might be more desirable than discrimination. Unfortunately, we currently lack the data to encorporate such features in our objective function. One study in honeybees looked at the output activation patterns of PNs in the AL and found that in these insects the physiological solution appears to be a correlation-based connectivity pattern (Linster et al., 2005). Apart from the anatomical differences between honeybees and fruit flies, this apparent contradiction could well be explained by such heterogeneity in the optimal odorant representation map. It is likely that different benchmark functions will lead to different optimal AL networks. Nevertheless the processing strategy characterized here can still be implemented whenever odors of similar ecological value need to be distinguished.

Our benchmark is based on physiological observations first studied in bees and later in several animal species, and considers concentration invariance (Sachse and Galizia, 2003). When dynamic odor-evoked activity patterns are represented in a multidimensional space, they form characteristic trajectories which quickly reach a set point, and then decay more slowly back to baseline (Friedrich and Laurent, 2001; Stopfer et al., 2003; Galán et al., 2004; Silbering et al., 2008). Interestingly, different concentrations of the same stimulus have transients that point into the same direction, but reach less far. Thus, in a static representation that considers only the set point, and not the dynamic trajectory (as done in this study), the vector direction contains the information about odor quality, and vector size contains the information about odor concentration (Sachse and Galizia, 2003).

Of course such a reduced description can not account for all aspects of information processing that have been shown to play an important role in the discrimination of odors in the AL of *Drosophila* and other species. Firstly, the transfer functions of individual PNs are not linear but saturate toward high activations. This is especially important for coding requirements like concentration invariance. Our linearity assumption is thus valid over a restricted dynamic range in which the transfer function evolves roughly linearly. Secondly, the neglect of temporal information rules out any sort of temporal coding strategy. Such strategies have been found to play an important role for example in locusts (Wehr and Laurent, 1996).

Our numerical simulations clearly indicate that a connectivity in the *Drosophila* AL that is based on input correlations is not superior in separating odor vectors as compared to randomly constructed networks, which show superior performance. Connectivity matrices that are based on the distance between glomeruli

fare similarly well as the random networks. This is an important result especially in the light of recent studies that characterize LN connections in *Drosophila* as homogeneous and global (Olsen and Wilson, 2008; Olsen et al., 2010). Here we show that under realistic input conditions derived from the DoOR database a globally inhibited network with homogeneous connection weights has the best odor separating performance. In neuroengineering scenarios in which neural networks are implemented on a VLSI chips this network is also the most simple to realize (Beyeler et al., 2010; Schmuker et al., 2011).

An analog picture is painted when looking at the total amount of inhibition that is exerted between the glomeruli. When measuring how much excess inhibition was produced by each network type we observed a less efficient inhibition only with the correlation-based network (**Figure 7**). This is important if one considers the metabolic costs of extensive inhibitory networks. This measure of metabolic costs and efficiency in our simplified model is admittedly far removed from a physiological description of metabolic efficiency and can only serve as a rough estimate of coding costs.

Lastly, we aimed to characterize the interaction between the statistical properties of the input matrix with the performance of our model networks. Our findings have important implications for the design of artificial olfaction systems. We have shown that different network types have contrasting odor separation performance depending on the response profiles of the odorant receptors. Global inhibition is most successful under conditions in which odor representations are evenly distributed over glomeruli, that is when the input network has high entropy (**Figure 8**). On the other hand, AL networks that inhibit glomeruli based on the degree of similarity between their response profiles (e.g., as calculated with a pair-wise correlation coefficient) perform better than random networks when the entropy of the input data is artificially reduced. This effect is somewhat surprising as in our simulations lower entropy networks are also less similar to the correlation matrix that underlies the connectivity. This effect can neither be attributed to the distribution of connection weights nor to the decorrelation between glomeruli because both the scrambled correlation-based network (with the same distribution) nor the flexible correlation-based network that was recomputed for each instance of the input (with higher decorrelation) perform worse than the simple correlation-based network.

Surprisingly, performance first increases slightly and then goes on to decrease drastically over all network types, as was to be expected. The initial increase is an expression of how much entropy is present in the original input data set. Some of the correlations in the dataset are decreased by the ordering of individual odor vectors when generating input matrices with lower entropy. If the receptor activations were mathematically optimally distributed across the whole input space this increase would disappear.

Summing up, our results show that a moderate level of global inhibition creates improved odor representations in the natural case of the *Drosophila* antennal lobe. Thus, this would be a good starting point for the processing of data from artificial chemosensory arrays. More importantly, however, we show that the best network is also dependent on the input statistics, which in turn is dictated both by the sensors used, and by the chemical to be detected. As a consequence, it is worth investigating the optimal network for each application. We propose a simplified simulating environment that allows for an efficient analysis of many possible networks, easy to be implemented when designing artificial noses.

REFERENCES

Asahina, K., Louis, M., Piccinotti, S., and Vosshall, L. B. (2009). A circuit supporting concentration-invariant odor perception in *Drosophila*. *J. Biol.* 8, 9.

Bargmann, C. I. (2006). Comparative chemosensation from receptors to ecology. *Nature* 444, 295–301.

Bazhenov, M., Stopfer, M., Rabinovich, M., Huerta, R., Abarbanel, H. D. I., Sejnowski, T. J., and Laurent, G. (2001). Model of transient oscillatory synchronization in the locust antennal lobe. *Neuron* 30, 553–567.

Beyeler, M., Stefanini, F., Proske, H., Galizia, C. G., and Chicca, E. (2010). "Exploring olfactory sensory networks: simulations and hardware emulation," in *Biomedical Circuits and Systems Conference BIOCAS, 2010 IEEE,* Paphos, 270–273.

Fiala, A., Spall, T., Diegelmann, S., Eisermann, B., Sachse, S., Devaud, J.-M., Buchner, E., and Galizia, C. G. (2002). Genetically expressed cameleon in *Drosophila melanogaster* is used to visualize olfactory information in projection neurons. *Curr. Biol.* 12, 1877–1884.

Friedrich, R. W., and Laurent, G. (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science* 291, 889–894.

Galán, R. F., Sachse, S., Galizia, C. G., and Herz, A. V. M. (2004). Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput.* 16, 999–1012.

Galizia, C. G., Münch, D., Strauch, M., Nissler, A., and Ma, S. (2010). Integrating heterogeneous odor response data into a common response model: A door to the complete olfactome. *Chem. Senses* 35, 551–563.

Gao, Q., Yuan, B., and Chess, A. (2000). Convergent projection of *Drosophila* olfactory neurons to specific glomeruli in the antennal lobe. *Nat. Neurosci.* 3, 780–785.

Hallem, E. A., and Carlson, J. R. (2006). Coding of odors by a receptor repertoire. *Cell* 125, 143–160.

Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117, 500–544.

Kazama, H., and Wilson, R. I. (2008). Homeostatic matching and nonlinear amplification at identified central synapses. *Neuron* 58, 401–413.

Laissue, P. P., Reiter, C., Hiesinger, P. R., Halter, S., Fischbach, K.-F., and Stocker, R. F. (1999). Three-dimensional reconstruction of the antennal lobe in *Drosophila melanogaster*. *J. Comp. Neurol.* 405, 543–552.

Linster, C., Sachse, S., and Galizia, C. G. (2005). Computational modeling suggests that response properties rather than spatial position determine connectivity between olfactory glomeruli. *J. Neurophysiol.* 93, 3410–3417.

Luo, S. X., Axel, R., and Abbott, L. F. (2010). Generating sparse and selective third-order responses in the olfactory system of the fly. *Proc. Natl. Acad. Sci. U.S.A.* 107, 10713–10718.

Olsen, S. R., Bhandawat, V., and Wilson, R. I. (2007). Excitatory interactions between olfactory processing channels in the *Drosophila* antennal lobe. *Neuron* 54, 89–103.

Olsen, S. R., Bhandawat, V., and Wilson, R. I. (2010). Divisive normalization in olfactory population codes. *Neuron* 66, 287–299.

Olsen, S. R., and Wilson, R. I. (2008). Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature* 452, 956–960.

Papadopoulou, M., Cassenaer, S., Nowotny, T., and Laurent, G. (2011). Normalization for sparse encoding of odors by a wide-field interneuron. *Science* 332, 721–725.

Sachse, S., and Galizia, C. G. (2002). Role of inhibition for temporal and spatial odor representation in olfactory output neurons: a calcium imaging study. *J. Neurophysiol.* 87, 1106–1117.

Sachse, S., and Galizia, C. G. (2003). The coding of odour-intensity in the honeybee antennal lobe: local computation optimizes odour representation. *Eur. J. Neurosci.* 18, 2119–2132.

Satoh, R., Oizumi, M., Kazama, H., and Okada, M. (2010). Mechanisms of maximum information preservation in the *Drosophila* antennal lobe. *PLoS ONE* 5, e10644. doi:10.1371/journal.pone.0010644

Schmuker, M., Yamagata, N., Nawrot, M. P., and Menzel, R. (2011). Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee. *Front. Neuroeng.* 4:17. doi: 10.3389/fneng.2011.00017

Schmuker, M., and Schneider, G. (2007). Processing and classification of chemical data inspired by insect olfaction. *Proc. Natl. Acad. Sci. U.S.A.* 104, 20285–20289.

Shang, Y., Claridge-Chang, A., Sjulson, L., Pypaert, M., and Miesenböck, G. (2007). Excitatory local circuits and their implications for olfactory processing in the fly antennal lobe. *Cell* 128, 601–612.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.* 27, 379–423.

Silbering, A. F., and Galizia, C. G. (2007). Processing of odor mixtures in the *Drosophila* antennal lobe reveals both global inhibition and glomerulus-specific interactions. *J. Neurosci.* 27, 11966–11977.

Silbering, A. F., Okada, R., Ito, K., and Galizia, C. G. (2008). Olfactory information processing in the *Drosophila* antennal lobe: anything goes? *J. Neurosci.* 28, 13075–13087.

Stocker, R. F. (1994). The organization of the chemosensory system in *Drosophila melanogaster*: a review. *Cell Tissue Res.* 275, 3–26.

Stocker, R. F., Lienhard, M. C., Borst, A., and Fischbach, K. F. (1990). Neuronal architecture of the antennal lobe in *Drosophila melanogaster*. *Cell Tissue Res.* 262, 9–34.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.

Vosshall, L. B., Amrein, H., Morozov, P. S., Rzhetsky, A., and Axel, R. (1999). A spatial map of olfactory receptor expression in the *Drosophila* antenna. *Cell* 96, 725–736.

Wang, J. W., Wong, A. M., Flores, J., Vosshall, L. B., and Axel, R. (2003). Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* 112, 271–282.

Wehr, M., and Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384, 162–166.

Wilson, R. I., and Laurent, G. (2005). Role of gabaergic inhibition in shaping odor-evoked spatiotemporal patterns in the *Drosophila* antennal lobe. *J. Neurosci.* 25, 9069–9079.

Wilson, R. I., Turner, G. C., and Laurent, G. (2004). Transformation of olfactory representations in the *Drosophila* antennal lobe. *Science* 303, 366–370.

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.

Received: 30 August 2011; paper pending published: 27 September 2011; accepted: 18 January 2012; published online: 08 February 2012.

*Citation: Proske JH, Wittmann M and Galizia CG (2012) Olfactory sensor processing in neural networks: lessons from modeling the fruit fly antennal lobe. Front. Neuroeng. **5**:2. doi: 10.3389/fneng.2012.00002*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 19 April 2012
doi: 10.3389/fneng.2012.00006

Non-linear blend coding in the moth antennal lobe emerges from random glomerular networks

Alberto Capurro[1]*, Fabiano Baroni[2], Shannon B. Olsson[3], Linda S. Kuebler[3], Salah Karout[1], Bill S. Hansson[3] and Timothy C. Pearce[1]

[1] *Department of Engineering, Centre for Bioengineering, University of Leicester, Leicester, UK*
[2] *Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, VIC, Australia*
[3] *Department of Evolutionary Neuroethology, Max-Planck-Institute for Chemical Ecology, Jena, Germany*

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Martin Nawrot, Freie Universität Berlin, Germany
Kerem Muezzinoglu, SAS Institute, USA

****Correspondence:***
Alberto Capurro, Department of Engineering, University of Leicester, Leicester LE1 7RH, UK.
e-mail: ac331@le.ac.uk; alberto.capurro@gmail.com

Neural responses to odor blends often exhibit non-linear interactions to blend components. The first olfactory processing center in insects, the antennal lobe (AL), exhibits a complex network connectivity. We attempt to determine if non-linear blend interactions can arise purely as a function of the AL network connectivity itself, without necessitating additional factors such as competitive ligand binding at the periphery or intrinsic cellular properties. To assess this, we compared blend interactions among responses from single neurons recorded intracellularly in the AL of the moth *Manduca sexta* with those generated using a population-based computational model constructed from the morphologically based connectivity pattern of projection neurons (PNs) and local interneurons (LNs) with randomized connection probabilities from which we excluded detailed intrinsic neuronal properties. The model accurately predicted most of the proportions of blend interaction types observed in the physiological data. Our simulations also indicate that input from LNs is important in establishing both the type of blend interaction and the nature of the neuronal response (excitation or inhibition) exhibited by AL neurons. For LNs, the only input that significantly impacted the blend interaction type was received from other LNs, while for PNs the input from olfactory sensory neurons and other PNs contributed agonistically with the LN input to shape the AL output. Our results demonstrate that non-linear blend interactions can be a natural consequence of AL connectivity, and highlight the importance of lateral inhibition as a key feature of blend coding to be addressed in future experimental and computational studies.

Keywords: olfaction, computational modeling, neural circuits, mixture processing, synaptic input, inhibitory interneurons, *Manduca sexta*

INTRODUCTION

Biological neural networks organize sensory inputs to produce a meaningful experience of the environment, but the way in which complex perceptual representations are produced by the olfactory system is not completely understood. For this reason, an understanding of the relation between the neural representation of a mixture and its single components constitutes an important problem in basic neuroscience. Moreover, such analyses can reveal general properties of perceptual representation in the nervous system, and derive principles that may be widely extended across species and sensory modalities (e.g., Rabinovich et al., 2008).

Natural odors are complex mixtures of different compounds. Within the olfactory system, the components of a blend often interact in a non-linear fashion within the olfactory system to affect the resultant neuronal response. This gives rise to so-called mixture interactions (Laing et al., 1989; Duchamp-Viret et al., 2003). The first interaction type is suppression, where the response to the mixture is less than at least one of the single components alone. A related category, in which the mixture evokes a response that is equivalent to the most effective single component, is known as hypoadditivity. The final interaction type is synergism, in which the mixture induces a response that is greater than the addition of the responses to the single components. A special case is linear addition, in which there is no interaction between the components, so the mixture evokes a response that is equal to the component sum.

The antennal lobe (AL) is the first synaptic relay in the olfactory pathway of insects, analogous to the mammalian olfactory bulb (for recent comparative review, see Martin et al., 2011). It consists of spheroidal bundles of neuropil known as glomeruli that contain synaptic contacts between receptors and second-order neurons. The axons of olfactory sensory neurons (OSNs) in the antennal nerve contact projection neurons (PNs) that constitute the output of the AL, and local interneurons (LNs) that communicate with other glomeruli. OSNs having the same type of receptor protein contact a specific glomerulus, giving rise to a spatial representation of chemical identity. In moths, approximately 250,000 OSNs from the antenna converge onto roughly 900 PNs and 360 LNs (Homberg et al., 1989; **Figure 1**). Interneurons in moths generally exhibit a broad symmetrical arborization pattern, contacting the majority of AL glomeruli (Manduca sexta: Matsumoto and Hildebrand, 1981; Kuebler et al., 2011; Reisenman et al., 2011). Although excitatory LNs have been found in *Drosophila* (Shang et al., 2007; Huang et al.,

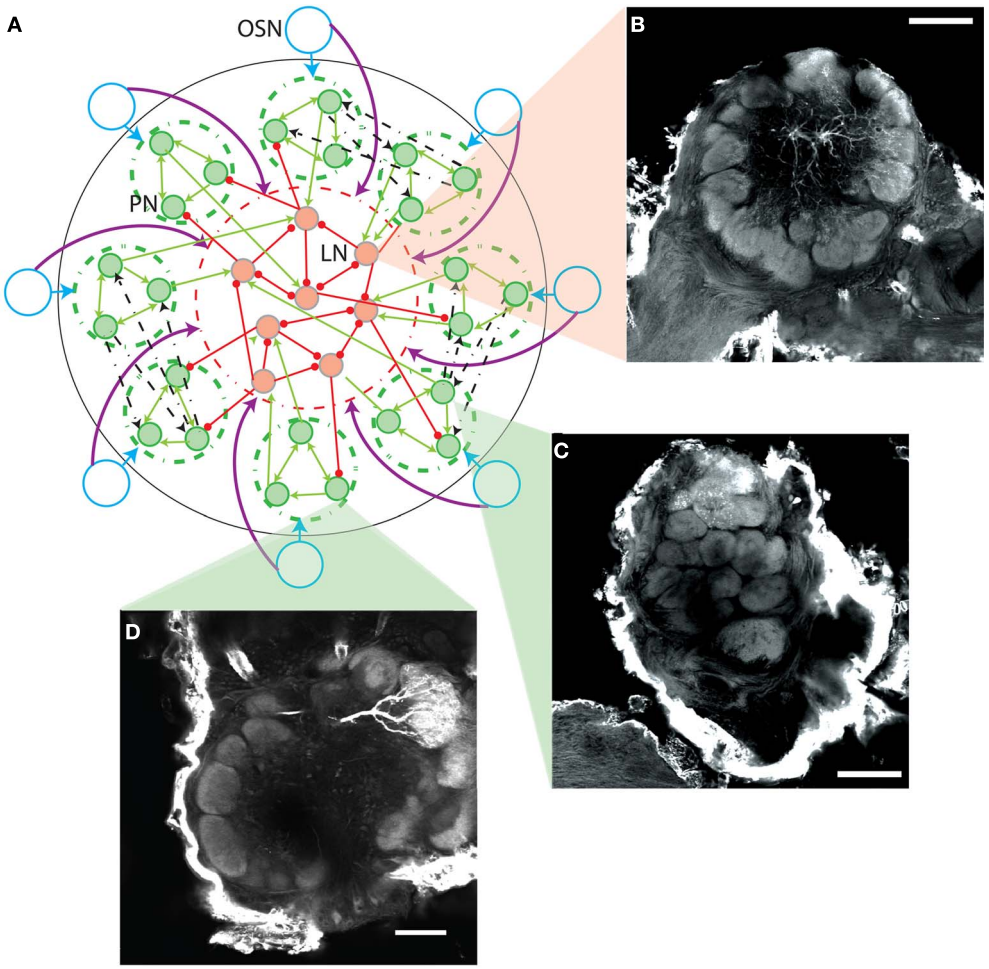

FIGURE 1 | (A) Scheme of AL model network connectivity. Each of the eight OSN types (blue circles) contact only PNs in their corresponding glomerulus (green dashed circles) as indicated by the blue arrows. OSN types contact every LN, as indicated by the purple arrows reaching the dashed red circle containing the LNs. The PNs have reciprocal excitation within the same glomerulus (green arrows within the dashed green circles), can contact the LNs (green arrows entering the dashed red circle), and have reciprocal multiglomerular excitatory contacts with another single glomerulus (dashed black arrows show a few examples). The LNs have reciprocal inhibition and can contact the PNs (red lines ending in circles). **(B–D)** Neuromorphic basis for the AL model. Confocal micrographs show three female *M. sexta* ALs with lines indicating their representative location in the model schema of **A**. Each panel extracts a single optical orthogonal slice. Neurobiotin-injected cells were stained with Alexa-conjugated Streptavidin. Pictures were obtained by confocal microscopy of three separate whole mount brain preparations using a 10×, 0.45-NA objective lens (C-Apochromat, Zeiss). Optical sections (1024 × 1024 pixel) were taken at intervals of 0.8 μm. **B** displays a lateral interneuron, while **C** and **D** show a multiglomerular and uniglomerular PN, respectively; scale bar: 100 μm.

2010), to date only inhibitory LNs have been located in moths (e.g., Reisenman et al., 2011). However, excitatory PNs that innervate two or more neighboring glomeruli have been identified in *Manduca sexta* hawkmoths (Homberg et al., 1988; Kuebler et al., 2011; **Figure 1**), and could provide a form of potential lateral excitation within the moth AL.

There exists considerable evidence indicating that the insect AL is not simply a relay station in the olfactory pathway, but constitutes the primary processing centre for blend information of the insect brain (e.g., Joerges et al., 1997; Linster and Smith, 1997; Hansson and Anton, 2000; Galizia and Menzel, 2001; Linster et al., 2005; Deisig et al., 2006, 2010; Carlsson et al., 2007; Silbering and Galizia, 2007; Krofczik et al., 2008; Lei and Vickers, 2008; Fernandez et al., 2009; Riffell et al., 2009a,b; Yamagata et al., 2009; Kuebler et al., 2011; Meyer and Galizia, 2011). Odor mixtures have been found to elicit mainly suppressive and hypoadditive responses within the insect AL, while examples of synergism are rare, as evidenced by both calcium imaging (e.g., Deisig et al., 2006, 2010; Carlsson et al., 2007; Silbering et al., 2008; Yamagata et al., 2009) and electrophysiological studies (e.g., Krofczik et al., 2008; Riffell et al., 2009a; Kuebler et al., 2011). In moths in particular, mixture interactions have been suggested to occur at the level of OSNs (e.g., Carlsson and Hansson, 2002; Hillier and Vickers, 2011) but are more commonly observed in second order neurons (e.g., Christensen et al., 1991; Lei and Vickers, 2008; Pinero et al., 2008).

In this study, we are interested in the network processing within the AL as a potential source of non-linear interactions between single components in the response to a blend. To assess this, we constructed a morphologically based computational model including populations of interconnected PNs and LNs using probabilistic rules. Parameters were calibrated to match the excitation/inhibition proportion observed in physiological recordings of AL neurons of the moth *M. sexta* (Kuebler et al., 2011). The model predicted most of the proportions of blend interactions found in the data, and the results are robust to changes in many parameters due to the underlying probabilistic connectivity. Finally, we used the model to investigate how different types of non-linear blend interactions could arise in terms of the synaptic input received by the PNs and LNs. We focused on the mean values of synaptic input during odor stimulus and control conditions, leaving the influence of the dynamic patterns of neural activity for a future study. Our results indicate that the array of blend interactions observed in the biological data can arise from network connectivity alone via sub-networks of inhibitory interneurons without requiring special intrinsic properties of the neurons themselves.

METHODS

NEURONAL RECORDINGS AND ODOR STIMULATION

We used intracellular recordings and morphological observations of AL neurons performed in the moth *M. sexta* (Lepidoptera, Sphingidae) by Kuebler et al. (2011), in addition to previous studies (see citations in "Network connectivity of the model") to construct a computational model of AL processing. The model parameters were calibrated to reproduce the proportion excitation/inhibition found in their physiological recordings (see "Network connectivity of the model"), and subsequently the proportions of blend interactions predicted by the model were independently compared with the recordings as a way of validating the model (**Figures 3** and **4**), which was in turn used to assess the role of synaptic input on the emergence of blend interactions (**Figures 5–8**). We include here a brief description of the recording methods (for full details see Kuebler et al., 2011).

Projection and interneurons were recorded intracellularly using sharp glass electrodes and stimulated for 500 ms at 10^{-4} dilution (in mineral oil) with (+) linalool, (−) linalool, phenyl acetataldehyde, benzaldehyde, hexanol, nonanal, or *trans*-2-hexenyl acetate (used instead of nonanal in some experiments), and *cis*-3-hexenyl acetate. Stimulus concentrations were equilibrated according to vapor pressure using a multicomponent stimulus device (Olsson et al., 2011). Neurons were presented with each of the seven odors separately, and the odors eliciting a response were tested together at the same concentration as a "blend." Active single components were finally tested separately at the total blend concentration. Mixture interactions were assessed as described below from response frequencies for each stimulus presentation, normalized to spontaneous activity as a ratio (Hz 1.5 s after stimulus onset/Hz 1.5 s before onset; including mechanical stimulus delay).

The complete data set consisted of 31 neurons tested each in one trial with the complete stimulus protocol. From this number, 20 neurons responded with excitation (including biphasic responses) and 11 responded with inhibition. Most of the recorded neurons (29 out of 31) could be classified as PNs or LNs by morphological analysis or by measuring the spike width, which nearly twice as large in LNs (for details see Kuebler et al., 2011). The cells that were morphologically labeled were 9 PNs and 5 LNs, while the criterion of the spike width allowed to classify 12 cells as PNs and 5 as LNs.

COMPUTATIONAL MODEL

Neuronal model

Individual neuron dynamics of LNs and PNs were modeled using a first-order differential equation (Chong et al., 2012) that described the evolution of the firing-rate activation variable of a neuron over time:

$$\tau \frac{da_i}{dt} = -a_i(t) + S\left(\sum_{j\in P} w_{i,j}a_j(t) + \sum_{k\in L} w_{i,k}a_k(t) + \sum_{d\in R} v_{i,d}r_d\right)$$

with $S(x) = x^3/(0.5^3 + x^3)$ for $x \geq 0$,

and $S(x) = 0$ for $x < 0$,

where a_i is the activation level of the i-th neuron, P is the subset of PN neurons, L is the subset of LN neurons, R is the subset of OSN neurons, $w_{i,j}$ is the strength of synaptic influence of j on the activity of i (similarly for $w_{i,k}$), $v_{i,d}$ is the strength of synaptic influence of the d-th OSN type on the activity of cell i-th, and r_d is the activity of the d-th OSN type. For PNs, $v_{i,d}$ is non-zero only for the connections coming from its corresponding OSN. S is a sigmoid function that limits the neuronal activity to values between 0 and 1. τ(10 ms for PNs and 20 ms for LNs) is the time constant of neuronal dynamics. We assumed that the larger dendritic ramifications of LNs can make them slower than the PNs, as has been done in previous articles that modeled the AL (e.g., Linster and Cleland, 2010; Chong et al., 2012).The neuronal activation function $S(x)$ (Chong et al., 2012) has a sigmoidal shape that accounts for saturation in activation level at high input values. We did not include noise added to each time step, but the initial values of a_i used in each realization were taken from a Gaussian random distribution with $\mu = 0.01$ and $\sigma = 0.0025$. The pre-stimulation control period started 200 ms after the onset of the simulation, when the system had already reached its resting state, so the transients due to the initial conditions (visible at times <50 ms in **Figure 2**) did not influence the results shown in **Figures 3–8**.

Network connectivity of the model

Figure 1A shows a general scheme of the network connectivity pattern. For simplicity only a few neurons of each type are represented, PNs with green filled circles and LNs with red filled circles. Excitatory synaptic contacts are represented with lines ending in arrow heads, and inhibitory contacts with lines ending in circles.

The model considers eight OSN types and eight glomeruli with 15 PNs each, making a total of 120 PNs. The total number of LNs was set to 40, resulting in a ratio PN/LN of three, which is realistic for the moth AL (Homberg et al., 1989). Thus, we are making a proportional reduction ($x = 8.75$) of the glomeruli number in the biological system of *M. sexta* (which is around 70; Grosse-Wilde et al., 2011), keeping a realistic number of PNs per glomerulus

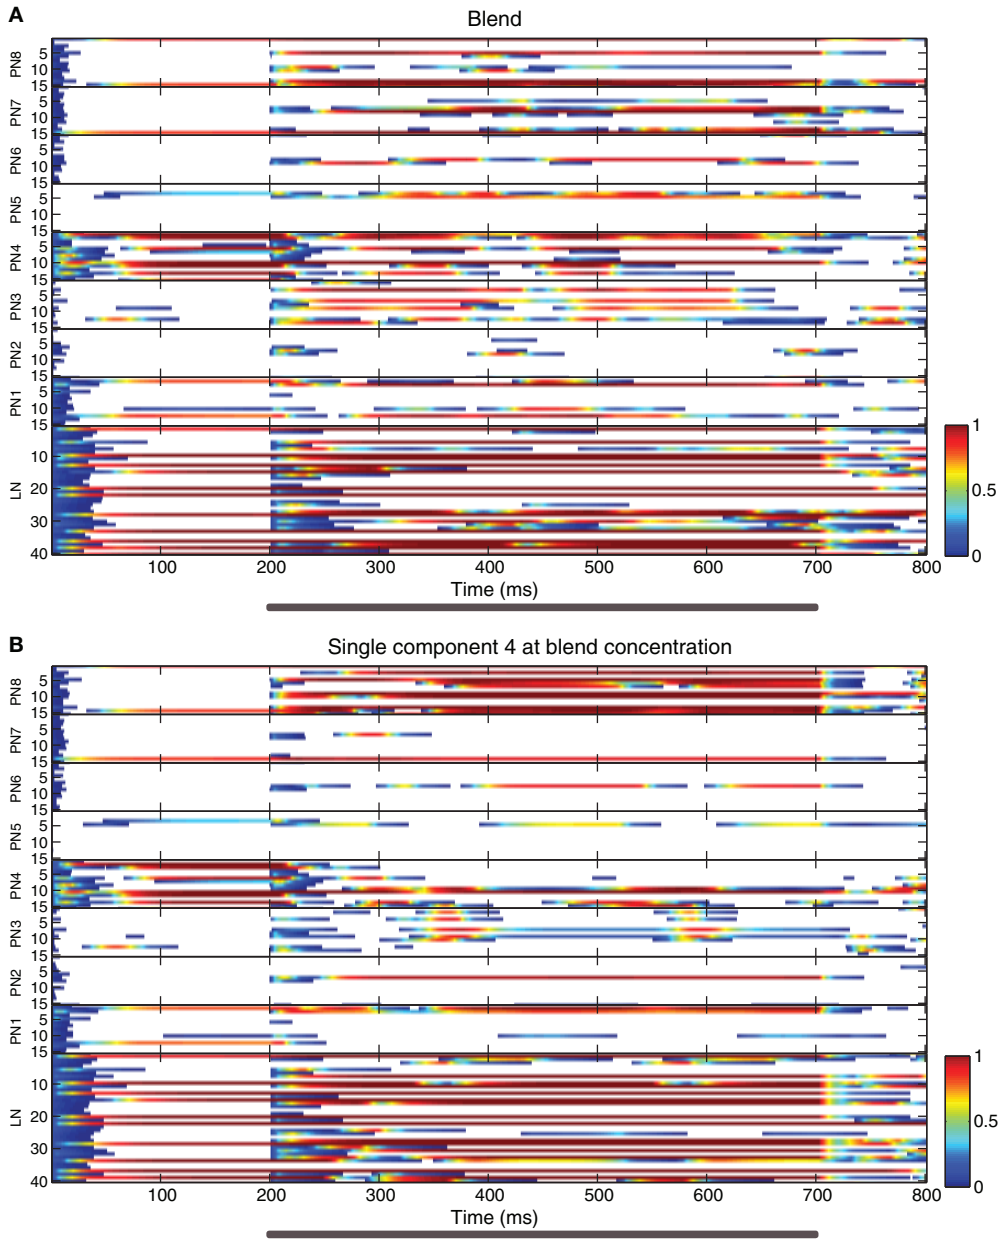

FIGURE 2 | Raster plots showing neuron activity in an example realization of the AL model. (A) Response to the homogeneous blend. **(B)** Response to a single component at blend concentration. The activity variable of each neuron is represented by a color code (shown in the scale bars at the lower right side of each panel) plotted against time. The ordinate's axis indicates the index numbers of the neurons: PNs belonging to each of the eight glomeruli (labeled PN 1 to PN 8, with 15 PNs per glomerulus) and 40 multiglomerular LNs (labeled LN). Odor stimulus duration is marked with gray bars under the abscissas.

and PN/LN proportions. This reduction resulted in a network scale in which the dynamic behavior of the model can be captured without an excessive increase in computational burden. We checked that the results regarding the proportions of blend interactions did not change significantly for a scaled-up version of the model having 64 glomeruli.

The connectivity between the different cell types is specified by its weight and connection probability. In each realization of the model, the values provided below for the weights were perturbed with Gaussian noise of mean zero and σ equal to 5% of the weight. OSNs of a given type are represented as a single unit (blue circles and arrows in **Figure 1A**) and only project to PNs of the corresponding glomerulus (dashed green circles) with a synaptic weight of 2.0. This configuration retains the generally dogmatic 1:1 principle between OSN type and glomerulus (Ressler et al., 1994; Mombaerts, 1996; Vosshall, 2000; Baker, 2008; Bruyne and Baker, 2008). Thus, the PNs of each glomerulus were activated only by their corresponding receptor type, as indicated by the blue

arrows. In contrast, the eight receptor types contacted every LN with a synaptic weight of 2.0, as indicated by the purple arrows of **Figure 1A**. Given the high convergence ratio between OSNs and other AL neurons within each glomerulus (300:1 for OSNs:PNs, ∼1000:1 for OSNs:LNs; Homberg et al., 1989), the connection probability of OSN to PN or LN contacts was set at 1.0.

The probabilistic rules used to establish the connectivity pattern of the network shown in **Figure 1A** are based on morphological studies of the *M. sexta* moth AL (e.g., Homberg et al., 1989), and also on direct observations performed by members of our group (**Figures 1B–D**). The location of the cells displayed in the morphological images in the context of the network scheme is indicated with triangular zooms. Here, we show the three major types of AL neurons recorded in the physiological study (Kuebler et al., 2011). As noted previously (Kuebler et al., 2011; Reisenman et al., 2011), LNs in the AL of the moth exhibited a broad, symmetrical arborization pattern throughout the AL (**Figure 1B**), interconnecting a large proportion of the glomeruli, as originally described (Matsumoto and Hildebrand, 1981). PNs exhibited two types of arborization patterns (Homberg et al., 1988). The majority of recorded PNs arborized in a single AL glomerulus, thus synapsing only with other PNs and LNs within that glomerulus (**Figure 1D**). However, a small proportion (2:8) of the stained PNs (Kuebler et al., 2011) arborized in two or more glomeruli (**Figure 1C**), thus synapsing with neurons within neighboring glomeruli.

To reflect these morphological observations in our model, PNs could have excitatory connections with other PNs of the same glomerulus with probability 0.8 and weight 0.37 (green arrows within the green dashed circles). In each glomerulus, we also randomly chose two PNs that could potentially contact PNs of another randomly chosen glomerulus with probability of 0.8 and weight 1.25 (the dashed black arrows in **Figure 1A** show only some examples). When a given glomerulus A received this type of contact from another glomerulus B, then B also innervated A with the same rule. In this way, each glomerulus was paired with another single glomerulus through this reciprocal excitatory connectivity (e.g., multi-glomerular PNs; **Figure 1C**). In the morphological data, we observed that multi-glomerular PNs generally innervated 2–3 neighboring glomeruli (Kuebler et al., 2011), however, our model does not include spatial representation. LNs could contact other LNs with a synaptic weight of −8.0 and probability of 0.25 (red lines within the dashed red circle of **Figure 1A**). In this way, both the PNs and the LNs had multiglomerular connectivity. LNs could contact PNs with a synaptic weight of −1.8 and a probability of 0.25, irrespectively of the glomerulus to which each PN belonged (red lines going beyond the dashed red circle in **Figure 1A**). PNs could feedback to LNs with a synaptic weight of 1.4 and probability of 0.15, as shown by the green arrows entering the dashed red circle in **Figure 1A**. This relatively sparse bidirectional connectivity between PNs and LNs ensured that not all PNs were inhibited during odor presentation, avoiding an exaggerated activation of the LN population at the same time (Chong et al., 2012). Synaptic weights of PN to PN and LN to LN connections were selected to reproduce the proportion of neurons responding with excitation/inhibition observed in the data (see below).

The values of synaptic weights and connection probabilities that we used in the model are biologically plausible. We selected them because they produced very sparse responses (involving less than 30% of the neuron population) with a proportion excitation/inhibition around 1.8:1, as observed in the physiological recordings (Kuebler et al., 2011). This ratio was particularly sensitive to the mean synaptic weight of lateral excitation and lateral inhibition. For uniglomerular PN to PN connection weight values of 0–0.3, we obtained ratios of around 2.5:1.0, while for larger weights the ratios decreased in a monotonic manner, reaching 1.0:1.0 for a weight of 0.6. In order to match the proportion observed in the recorded data we used a weight of 0.37, as indicated above. For multi-glomerular PN to PN connections, the proportion also decreased with an increase in weight, varying from 2.4:1.0 to 1.26:1.0 for weights of 1.0–1.5, so we chose the intermediate weight of 1.25 that matched the recordings. In the case of the LN to LN connections the proportion excitation/inhibition showed a U-shaped profile for a weight range from −4.0 (4.0:1.0) to −10.0 (2.2:1.0) with a minimum in −60 (1.5:1), and so we selected a weight of −8.0 that matched the ratio observed in the data (1.8:1.0). We did not optimize the parameters to match the proportions of blend interactions being the main focus of our study, which were indeed very robust to changes in the lateral excitation and to increases in the lateral inhibition. Some parameters were selected following a previous study (Chong et al., 2012), such as the probability of LN to LN contacts and uniglomerular PN to PN contacts, as well as the use of low levels of bidirectional coupling between PNs and LNs. Note that reliable estimates for these parameters are not available in the moth. Furthermore, the high dimensionality of our model prevents an exhaustive parameter search. Hence, even if our parameters are within a physiological range, we cannot rule out the existence of another physiologically plausible parameter set that also results in the 1.8:1 ratio between excited and inhibited responses.

Odor stimulus in the model

In line with the blends used in the physiological experiments, we considered odor stimuli to exist within a five-dimensional space ($Q = 5$), where each coordinate represents the concentration, c_h, of a single chemical component, h. The network was stimulated independently with the single component odors, the homogeneous blend and the single component odors at blend concentration. Single component odorants had zero concentration for each odor dimension but one, such that the five separate stimulus vectors ($q = 1, \ldots, Q$)

$$\overrightarrow{c_q} = [c_{q1}, \ldots, c_{qh}, \ldots, c_{qQ}],$$

contain the components $c_{qh} = 0$ for $h \neq j$ and $c_{hq} = 1$ for $h = q$. Thus, all component concentrations were normalised to one and dimensionless for simplicity. The homogeneous blend stimulus was defined as all components combined simultaneously, $c_h = 1 \,\forall\, h$, and the single component odors at blend concentrations were derived from the single component odors $\overrightarrow{c_{q(\text{blend})}} = 5\overrightarrow{c_q}$. Thus, in total there are eleven concentration vectors, five single component odors, $\overrightarrow{c_q}$, the homogeneous blend concentration, $\overrightarrow{c_{\text{blend}}}$, and the single component odors at blend concentration

$\overrightarrow{c_{q(\text{blend})}}$. The range of input concentrations used in the simulations (0–5) was selected to match the range used in the experiments.

OSN activation in response to stimulus presentation was calculated as follows. First, each OSN type odor dimension pair was assigned a binding value, a (analogous to affinity to quantify the strength of binding), drawn from a Gaussian distribution with mean 0.5 and standard deviation 0.1. These values characterise the binding strength of the different pure chemicals for each of the molecular receptor types. Then, for a given OSN type, the different input concentration vectors, $\overrightarrow{c}$, were multiplied component by component with the binding values $\overrightarrow{a} = [a_1,\ a_2, \ldots,\ a_5]$ to obtain the binding vector $\overrightarrow{z} = [c_1 a_1,\ c_2 a_2, \ldots,\ c_5 a_5]$. This scheme provided a plausible description of the typical binding properties found in insect olfaction systems (Hallem and Carlson, 2006), since it captures the common observation that generalist molecular receptors have variable affinity for many chemicals, and cases of very high or very low affinity are relatively rare among general odors.

For a given odor presented to the model, the eight binding vectors corresponding to each OSN type were merged into a matrix z_{dq} of 40 numbers (five columns and eight rows), each specifying the binding of a given compound q (columns) to each of the OSN types d (rows). Each binding value z_{dq} was then passed through a sigmoid function that represented the dose-response curve of OSN activity to a given component

$$r_{dq} = \frac{\lambda}{1 + e^{-\alpha(z_{dq}-\gamma)}} + \eta$$

to obtain a matrix r_{dq} which specifies the activities of each OSN type due to each component. The parameters α, γ, η, and λ control the slope, horizontal shift, vertical shift, and amplitude of the sigmoid, respectively. For each element of the match in order to create a diversity of tunings, parameter values were drawn from uniform distributions in the range [0,5], [0,4], [0,0.1], and [0,1], respectively. In this way, we obtained 40 different sigmoid functions (corresponding to each element of the matrix) that were used to calculate the output activity of each receptor type d due to the binding of each chemical q. Receptor activations for each OSN type d were obtained by summing the columns of the matrix r_{dq}. The sigmoid functions were different for different odorants across the receptor types (Hallem and Carlson, 2006).

Since we did not want to generate non-linearities in the OSN response to blends at the periphery level, we assumed a linear summation in OSN activities of the blend components. Thus, the columns of this matrix were summed linearly $r_d = \sum_q r_{dq}$ to obtain an eight-dimensional vector with the total activity of each receptor type which was the actual input from the OSNs to the PNs and LNs of the model (see Equation 1). It has been reported by Rospars et al. (2008), that certain type of competitive scheme may occur in around half of the receptors, while others show more complex allosteric interactions. However, since many details of the periphery function are still unknown in the moth, we did not want to add further hypothetical non-linear behaviors at the receptor level of our model, because they would obscure our assessment of the interaction types at the neuronal and network levels. In addition, this generalization makes the model more generally applicable for other non-moth systems.

When a stimulus was present, we added a small positive offset to OSN activations. This offset (set to 1.0) represents the recruitment of non-specific OSNs which activate at very low concentration and was assumed to be equal for each glomerulus for simplicity. For no binding at the input, the values of d were small positive numbers, simulating some degree of spontaneous discharge in the OSNs (Hallem et al., 2004). Notice that we are modeling the receptor activity at a population level and the magnitudes are dimensionless. The parameter ranges were adjusted qualitatively to obtain a family of sigmoid functions that produced similar dose-response curves in terms of dynamic range (up to 3 log units), sensitivity, and response intensity axes to those found experimentally using optical recordings (Carlsson and Hansson, 2003). This scheme resulted in very little blend interactions at the OSN level (see "High correlation in blend interaction types between physiological data and computational model"), which is in agreement with recent observations performed using optical recordings (Kuebler et al., 2012).

PROCEDURE TO COMPUTE THE BLEND INTERACTIONS

The procedure used to determine the proportions of blend interactions in the response of single neurons was the same for the model simulations and the AL intracellular recordings, using only the neurons that responded to the stimulus. The stimulation was performed with the single components, the homogeneous blend, and the single components at blend concentration, using a stimulus pulse that lasted for 500 ms. Responsive neurons were then classified according to the relationship between the blend response and the responses to single components. Briefly, the maximum and the standard deviation of the responses to single components ($\max_S, \sigma_S$) and to single components at blend concentration ($\max_{SB}, \sigma_{SB}$) were computed. Then, classification was performed as follows: (1) suppression: blend response $< \max_S - \sigma_S$, (2) hypoadditivity: $\max_S - \sigma_S <$ blend response $< \max_S + \sigma_S$, (3) linear addition: $\max_S + \sigma_S <$ blend response $< \max(\max_S + \sigma_S, \max_{SB} + \sigma_{SB})$, and (d) synergy: blend response $> \max(\max_S + \sigma_S, \max_{SB} + \sigma_{SB})$.

The index of cell activity used to quantify the response in the recordings was the mean firing rate. In the case of the model we used the mean of the activity variable (see "Neuronal model") which ranges from 0 to 1. In both cases, we compared the activity evoked by the stimulus with the activity of a control period before the stimulus onset. The response was taken as the difference between the activities in these two time windows. In the AL of the moth, some neurons exhibited biphasic responses, which consist of excitation followed by inhibition. Our model does not account for this response type, and we thus pooled neurons that responded with excitation with neurons that showed biphasic responses, since the latter also consist of a net increase in firing rate. Hence, we concentrated here on the mean values of the neuronal activity leaving the assessment of the dynamic patterns to be presented in a future study.

As explained above, we followed parallel procedures to calculate the proportions of blend interactions in the recordings and

the simulations. In both cases, we assessed mean response intensity only, and not response duration. This is important because in the physiological data the response outlasted the stimulus, while in the model neurons are endowed with no intrinsic properties and thus the response ended with the stimulus offset. In order to account for the entire response period, the response of the model was considered to be the difference between the mean activity value in the 500 ms following and preceding the stimulus, while in the intracellular recordings the odorant response was assessed for 1400 ms following and preceding the stimulus onset. An assessment of the model with a stimulus lasting 1400 ms gave the same results as using 500 ms. In the model, a neuron was considered to be responsive when the difference in mean activity between the stimulus (with the blend or a single component at low concentration) and the control period exceeded a value of 0.1, which corresponds to 10% of the maximum activity value. This gives a response threshold that is close to one observed in the physiological data (Kuebler et al., 2011).

RESULTS

HIGH CORRELATION IN BLEND INTERACTION TYPES BETWEEN PHYSIOLOGICAL DATA AND COMPUTATIONAL MODEL

We performed 100 realizations of the computational model (network connectivity shown in **Figure 1**), each with different receptor binding matrices and sigmoidal functions for the model OSNs. **Figure 2** displays raster plots of an example realization showing the activity of all cells during the stimulus with the homogeneous blend (panel A) and with a single odorant at the blend concentration (panel B).

The biological and computational comparisons of blend interactions were performed by selecting a set of synaptic weights and connection probabilities that allowed a biologically reflective sparsity in the AL neuronal response including an approximate ratio of 1.8:1.0 between neurons that responded with excitation and inhibition, as observed in the recordings (Kuebler et al., 2011). In both cases, excitation was more prevalent in LNs, while PNs exhibited similar levels of excitatory or inhibitory responses. For PNs, the proportion excitation/inhibition was 0.86:1.0 in the model and 1.3:1.0 in the recordings where 12 cells responded with excitation and nine with inhibition. For LNs, this proportion was 3.0:1.0 in the model and also 3.0:1.0 in the recordings where six LNs responded with excitation and two with inhibition.

Responses were classified into four types according to the activity of each neuron in response to the blend vs. its single components. When there is no interaction between the components, the response evoked by the blend is equal to the linear sum of the responses to the individual components. In the case of synergism, the response exceeds the linear sum. Suppression is a reduced response with respect to the single components, while in the case of hypoadditivity (often referred to as overshadowing or blocking) at least one of the components of the blend is ignored and the response resembles that of the most effective component.

The proportions of blend interaction types are shown in **Figure 3** for neurons that respond with excitation (**Figure 3A**) and inhibition (**Figure 3B**). The blue columns represent the model results, while the red squares represent the values found in the experimental recordings of AL neurons. For both the recordings and the simulations, we found that most interactions between odorants were non-linear, and within the non-linear interactions hypoadditivity (blend response = single components) and suppression (blend response < single components) were more common than synergism (blend response > single components at blend concentrations). For the responses consisting of excitation (**Figure 3A**), the red squares are within the error bars of the blue columns, close to its mean value. This indicates a good agreement between the physiological data and the model. This is also the case for synergism and linear addition in the

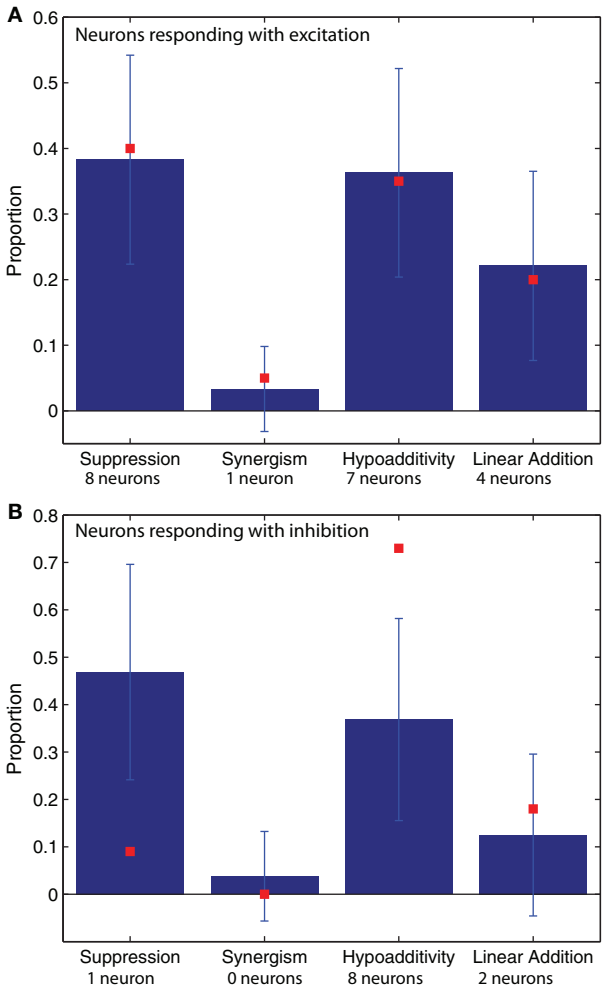

FIGURE 3 | The proportions of blend interaction types found in 100 model realizations ($\mu \pm$ SD) are shown with blue bars. The data values obtained from the experimental recordings ($N = 31$ neurons) are plotted with red squares. **(A)** Neurons that respond with excitation ($N = 20$). **(B)** Neurons that respond with inhibition ($N = 11$). The number of neurons used to compute the proportions represented by the red squares are provided under their corresponding abscissa labels. For the model simulations of neurons that respond with excitation (blue bars in **A**) the neuron numbers were 901 for suppression, 822 for synergy, 546 for hypoadditivity, and 81 for linear addition. For the neurons that respond with inhibition (blue bars in **B**) the neuron numbers were 622 for suppression, 520 for synergy, 179 for hypoadditivity, and 54 for linear addition.

neurons that respond with inhibition (**Figure 3B**). However, in this case the model does not fit the experimental results for suppression and hypoadditivity but exhibits more of the former and less of the latter type of blend interaction than the physiology (**Figure 3B**). These two categories are both forms of suppression, and the sum of their relative observations (about 0.8) is equal between the recordings and the model (**Figure 3B**). We suspect that intrinsic neuronal properties that were not included in the model are responsible for this discrepancy, altering the response of the neurons in the vicinity of the threshold between both interaction types.

The results shown in **Figure 3** did not change when we increased or decreased the weight of the lateral excitation or when we increased the weight of the lateral inhibition in the ranges specified in section "Network connectivity of the model". For weaker lateral inhibition values, we observed an increase in linear addition accompanied with a decrease of the same magnitude ($\sim$0.1) in suppression in the neurons that respond with excitation. Conversely, the proportion excitation/inhibition was very sensitive to variations in the values of these parameters (see "Network connectivity of the model").

To assess whether the observed proportions of blend interaction types arose from the AL network (PNs and LNs) or from the OSN periphery of the model, we ran a separate simulation using a reduced version of the model. In this simulation, we suppressed all inhibitory and excitatory connections within the AL and the input/output function of PNs and LNs was linear [$S(x) = x$], so the output of an AL neuron was simply a linear function of the input received from the OSN periphery. Under these conditions, we observed that all responses consisted in excitation, resulting in linear addition for most AL neurons (for every LN and for ¾ of PNs), while the remaining cases showed hypoadditivity. This is in close agreement with experimental observations found with optical recordings of the compound input (Kuebler et al., 2012) and suggests that the periphery has little influence on AL blend interactions in the moth. The same result was found when we suppressed only the lateral inhibition (LN to LN and LN to PN connections) using $S(x) = x/10$ to avoid an explosion of excitation in the simulation. This means that the result depicted in the blue bars of **Figure 3** arises principally from the AL network with a fundamental contribution made by the lateral inhibition. We also performed a control simulation with the full AL connectivity but using the linear input/output function $S(x) = x$ (with negative values rectified to 0) and observed that the results of **Figure 3** and the ratio excitation/inhibition were not significantly modified, implying that they are not due to the particular form of the non-linear squashing function used for AL neurons (see "Neuronal model").

As explained in "Neuronal recordings and odor stimulation," most of the recorded neurons (29 out of 31) could be classified as PNs or LNs (Kuebler et al., 2011). Using this classification, we compared the proportions of blend interaction types for PNs and LNs of different response types (**Figure 4**). Neurons that responded with excitation are depicted in panels A (PNs) and B (LNs), while neurons that responded with inhibition are depicted in panel C (PNs) and D (LNs). Despite the low cell numbers for the physiological measurements, we found that the model follows the trends observed in the recoded data for both PNs and LNs that responded with excitation (**Figures 4A,B**). In the case of PNs that respond with inhibition (**Figure 4C**) the model matched the proportions of the data only for synergism and linear addition. Comparing **Figures 3** and **4** we can conclude that the differences found in the pooled data set (**Figure 3**) are still apparent when assessing PNs and LNs separately (**Figure 4**).

In this work, as in previous experimental studies (Duchamp-Viret et al., 2003; Silbering and Galizia, 2007; Kuebler et al., 2011), blend interaction types (with the exception of synergy) were defined with respect to single component odors delivered at a concentration equal to the blend concentration divided by the number of single components tested (e.g., we used five components with blend concentration of five, resulting in a single concentration of 1). We also confirmed that the results presented in **Figure 3** are similar for blends of 2–7 components, or for blend concentrations from 1.25 to 10. These alterations did, however, affect the proportion or excitation/inhibition, which could be compensated by altering the levels of lateral inhibition. Thus, although our model results in monotonous responses to changes in input concentration, more experimental data is necessary to evaluate the generalization of our simulations over a wide range of input concentrations.

In this subsection, we have shown that our computational model, with a parameter set calibrated in order to reproduce the proportion of excitation/inhibition and response sparsity found in the physiological recordings, can also reproduce the proportions of blend interactions types without any additional tuning. While the model matches the data remarkably well for neurons that respond with excitation, hypoadditivity was underestimated and suppression was overestimated in neurons that responded with inhibition, although the sum of their relative observations was matched. Under these conditions of biologically plausible synaptic interactions, we assess in the following how blend coding is shaped by the different sources of synaptic input received by individual AL neurons.

LN INPUT SHAPES BLEND INTERACTIONS AND RESPONSE TYPES IN AL

Our goal was to understand how the different types of responses and blend interactions emerged in the AL network. Hence, we analyzed the synaptic inputs to AL neurons and their changes as the stimulus was varied. In **Figure 5A** we plot the difference in total synaptic input between the blend stimulation period and the pre-stimulation period, multiplied by the corresponding synaptic weights for different response categories in the model. This magnitude is referred to as synaptic input. As expected, neurons that responded with excitation experienced an increase in net input with blend stimulation, while neurons that responded with inhibition showed a decrease in their net input (**Figure 5A**).

Figures 5B and **C** display the synaptic input arriving from the LNs (left panels), PNs (middle panels), and OSNs (right panels) to the PNs (**Figure 5B**), and LNs (**Figure 5C**). For both cell types, changes in lateral inhibition were the main determinants of response type. In the case of PNs, however, other PNs and OSNs also contributed to the response. In the case of LNs (**Figure 5C**) the only relevant input arrived from other LNs, with

no significant contribution carried by the input arriving from PNs and OSNs (middle and right panels in **Figure 5C**).

As interaction types are defined by the relation between the blend response and the responses to the single components, we performed a similar analysis on the synaptic inputs to each neuron considering the difference between the input during blend stimulation and the mean input during single component presentations. This magnitude, referred as change in synaptic input,

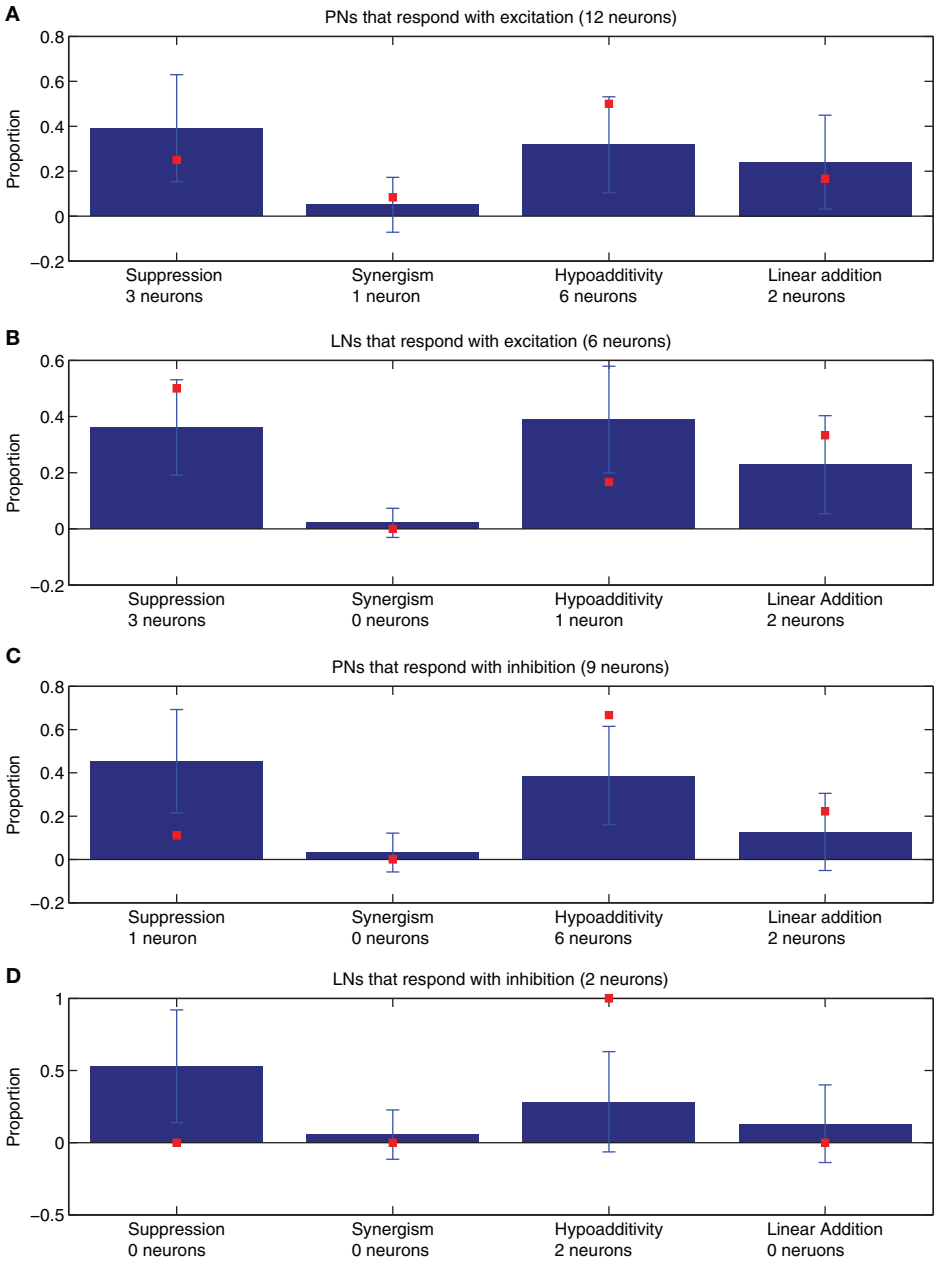

FIGURE 4 | The proportions of blend interaction types found in 100 model realizations (μ ± SD, blue bars) are compared with the data values obtained from the experimental recordings of the neurons that could be successfully classified as PNs or LNs (*N* = 29 neurons, red squares). (A) PNs that respond with excitation ($N = 12$). **(B)** LNs that respond with excitation ($N = 6$). **(C)** PNs that respond with inhibition ($N = 9$). **(D)** LNs that respond with inhibition ($N = 2$). The number of neurons used to compute the proportions represented by the red squares are provided under their corresponding abscissa labels. For the model simulations of PNs that respond with excitation (blue bars in A) the neuron numbers were 397 for suppression, 306 for synergy, 240 for hypoadditivity, and 47 for linear addition. For the LNs that respond with excitation (blue bars in B) the neuron numbers were 504 for suppression, 516 for synergy, 306 for hypoadditivity, and 34 for linear addition. For the PNs that respond with inhibition (blue bars in C) the neuron numbers were 506 for suppression, 457 for synergy, 151 for hypoadditivity, and 42 for linear addition. For the LNs that respond with inhibition (blue bars in D) the neuron numbers were 116 for suppression, 63 for synergy, 28 for hypoadditivity, and 12 for linear addition.

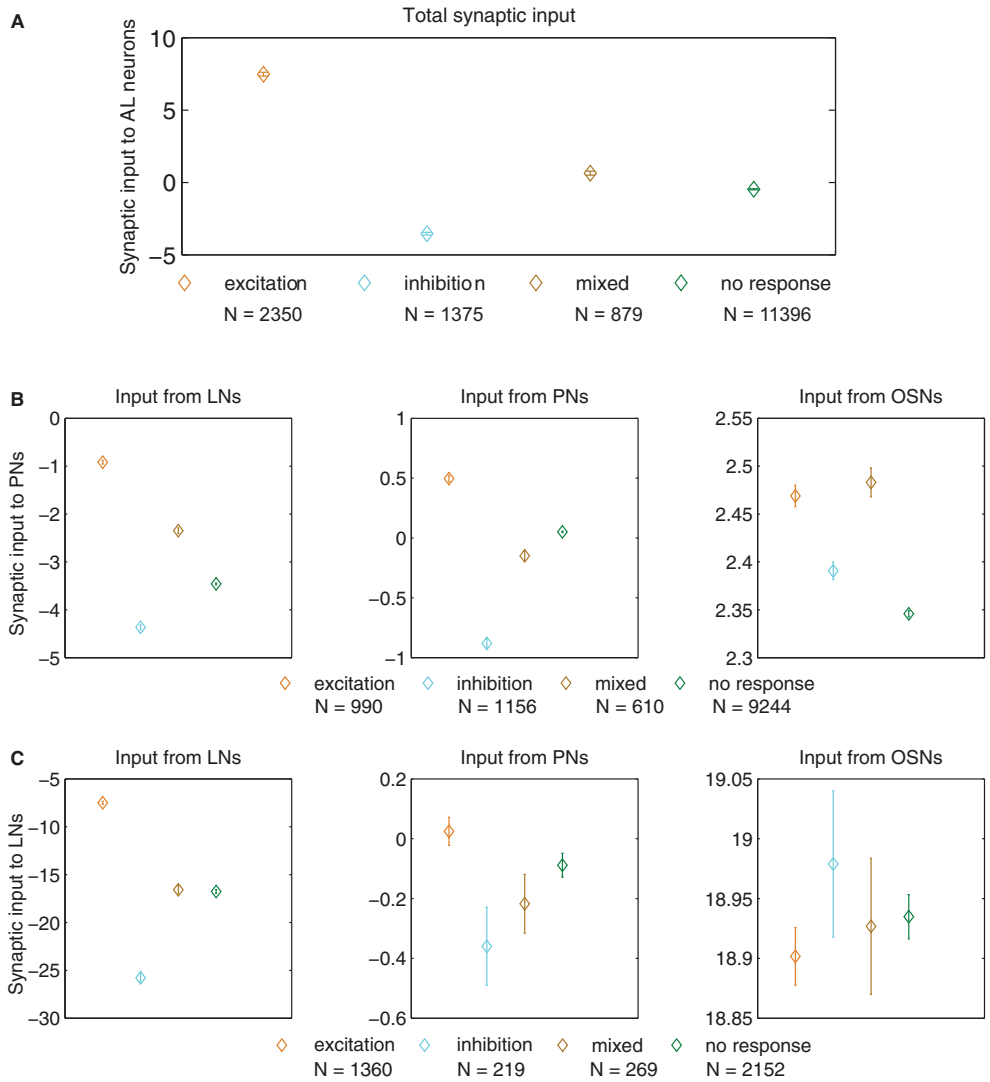

FIGURE 5 | Synaptic input (defined as the difference in input between the blend stimulation period and the pre-stimulation control period, multiplied by the corresponding synaptic weights) for neurons of different response type ($\mu \pm$ SEM for 100 model realizations). (A) Total input to AL neurons. All comparisons reached $p < 1 \times 10^{-6}$ in the *t*-test after Bonferroni correction, the most important result being responses consisting of excitation vs. inhibition (orange and cyan diamonds). In panels **B** and **C**, we show the input arriving from different types of neurons (LNs, PNs, and OSNs) to AL neurons (PNs and LNs). **(B)** Input to PNs arriving from LNs (left panel), PNs (middle panel), and OSNs (right panel). **(C)** Input to LNs arriving from LNs (left panel), PNs (middle panel), and OSNs (right panel). The *p* values for cases in which we found significant differences in the *t*-test ($p \leq 0.01$, after Bonferroni correction) are listed in the following, the most important result being responses consisting of excitation vs. inhibition (orange and cyan diamonds). For the input from LNs to PNs (left panel of **B**) all comparisons reached $p < 1 \times 10^{-6}$. For the input from PNs to PNs (middle panel of **B**) comparisons reached $p < 1 \times 10^{-6}$. For the input from OSNs to PNs (right panel in **B**) excitation vs. inhibition $p < 1 \times 10^{-6}$, excitation vs. no response $p < 1 \times 10^{-6}$, inhibition vs. mixed response $p < 1 \times 10^{-6}$, inhibition vs. no response $p < 2 \times 10^{-5}$, and mixed vs. no response $p < 1 \times 10^{-6}$. For the input from LNs to LNs (left panel in C) all comparisons reached $p < 1 \times 10^{-6}$, with the only exception of mixed vs. no response that was not significant.

was calculated as the difference between the blend and the average single components at low concentration, multiplied by the corresponding synaptic weights (**Figures 6–8**). In **Figure 6** we depict the total change in synaptic input for neurons that respond with excitation (**Figure 6A**) and inhibition (**Figure 6B**), irrespective of their neuronal type (PNs or LNs). The following two figures unfold this result separating PNs from LNs and specifying the different sources of input. In **Figure 7** we show neurons that respond with excitation (PNs in **Figure 7A** and LNs in **Figure 7B**), while in **Figure 8** we show neurons that respond with inhibition (PNs in **Figure 8A** and LNs in **Figure 8B**).

Figure 6 shows that the change in total input grew in the order: suppression, hypoadditivity, and linear addition in neurons that responded with excitation (**Figure 6A**), and it decreased following the same sequence in the case of neurons that responded with inhibition (**Figure 6B**).

Figure 7 shows that the input change arriving from LNs to LNs that responded with excitation (left panel in **Figure 7B**) grew

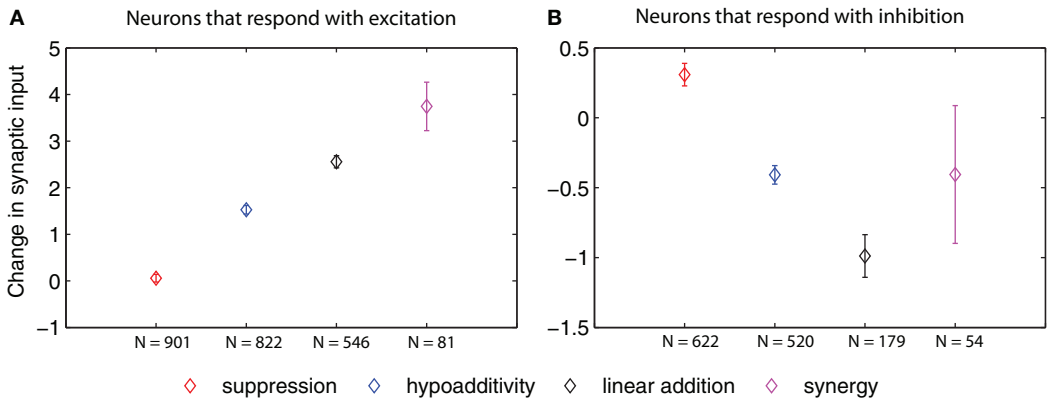

FIGURE 6 | Change in synaptic input (defined as the difference in input between the blend response and the average response to the single components at low concentration, multiplied by the corresponding synaptic weights) for AL neurons of different blend interaction types ($\mu \pm$ SEM for 100 model realizations). (A) Total change in synaptic input arriving to neurons that respond with excitation. **(B)** Total change in synaptic input arriving to neurons that respond with inhibition. The p values for cases in which we found significant differences in the t-test ($p \leq 0.01$, after Bonferroni correction) are listed in the following. In panel **A**: suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p < 1 \times 10^{-6}$; suppression vs. synergy $p < 1 \times 10^{-6}$; hypoadditivity vs. linear addition $p =< 1 \times 10^{-6}$, hypoadditivity vs. synergy $p < 1 \times 10^{-6}$. In panel **B**: suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p < 1 \times 10^{-6}$; hypoadditivity vs. linear addition $p = 3.8 \times 10^{-4}$.

monotonically from suppression to synergy, creating the image of an ascending ladder. A similar picture applies for the input change arriving from the LNs to the PNs that responded with excitation (left panel in **Figure 7A**), although in this case linear addition and synergy (the last step of the ladder) were not significantly different. This suggests that a stronger lateral inhibition in response to the blend established suppression, while small changes and decreases in lateral inhibition resulted in hypoadditivity and linear addition (or synergy), respectively. The change in input arriving from PNs to PNs that responded with excitation (middle panel in **Figure 7A**) again created the trend of an ascending ladder with a missing last step. This is not the case for the change in input arriving from PNs to LNs that responded with excitation (middle panel in **Figure 7B**) where no significant differences were found. Regarding the input change from OSNs, a decrease for PNs that responded with excitation (right panel in **Figure 7A**) established suppression, but there were no significant differences in input changes from OSNs to LNs that responded with excitation for any interaction type (right panel in **Figure 7B**).

Figure 8 shows that the change in input establishing suppression for responses consisting of inhibition likewise tended to differ from that establishing the other interactions types. For PNs that responded with inhibition (left panel in **Figure 8A**) the input change from LNs in establishing suppression was greater than both hypoadditivity or linear addition. This situation repeated for the input from OSNs to PNs that responded with inhibition (right panel in **Figure 8A**). For the input from PNs to PNs that responded with inhibition, the input change creating suppression was greater than in the other three categories (middle panel in **Figure 8A**). For the input from LNs to LNs that responded with inhibition the input change creating suppression was greater than linear addition (left panel in **Figure 8B**), but no significant differences were found in the input changes from PNs and OSNs to LNs (middle and right panels in **Figure 8B**). The increased change in synaptic input that created suppression for responses consisting of inhibition resulted from both a decreased inhibition coming from the LNs, and a stronger excitation coming from the OSNs and PNs.

In summary, **Figures 7** and **8** indicate that blend interactions were shaped mainly by the input changes coming from LNs in all AL cells of the model, with an agonistic contribution of smaller magnitude from PNs and OSNs in the case of PNs.

DISCUSSION

COMPARISON OF BLEND INTERACTIONS IN THE MODEL AND RECORDINGS

The proportions of non-linear blend interaction types observed in the biological data were accurately predicted by the simulation results in most cases. The matching is almost perfect for neurons that respond with excitation. For neurons that respond with inhibition, a good agreement was found for synergism and linear addition, while in the case of suppression and hypoadditivity the model could only match the sum of both categories, but not the actual proportions of each. Both hypoadditivity and suppression are forms of suppressive interactions (Kuebler et al., 2011), hence this discrepancy is quantitative rather than qualitative. Future research can elucidate whether more realistic single-cell properties, or the presence of inhibitory responses at the OSN level (e.g., Hallem and Carlson, 2006) could result in a better match. Indeed, the high correlation between the model and the physiological data is surprising considering the simplicity of the model employed, and must therefore originate in the morphologically based pattern of neuronal population interactions. As stated in section "High correlation in blend interaction types between physiological data and computational model", the neuronal activation function $S(x)$ (Chong et al., 2012) is not responsible for the proportions of blend interactions observed in **Figure 3**, which are also found to hold over a wide dynamic range of stimulus. This type of non-linearity is widely used in the field of computational

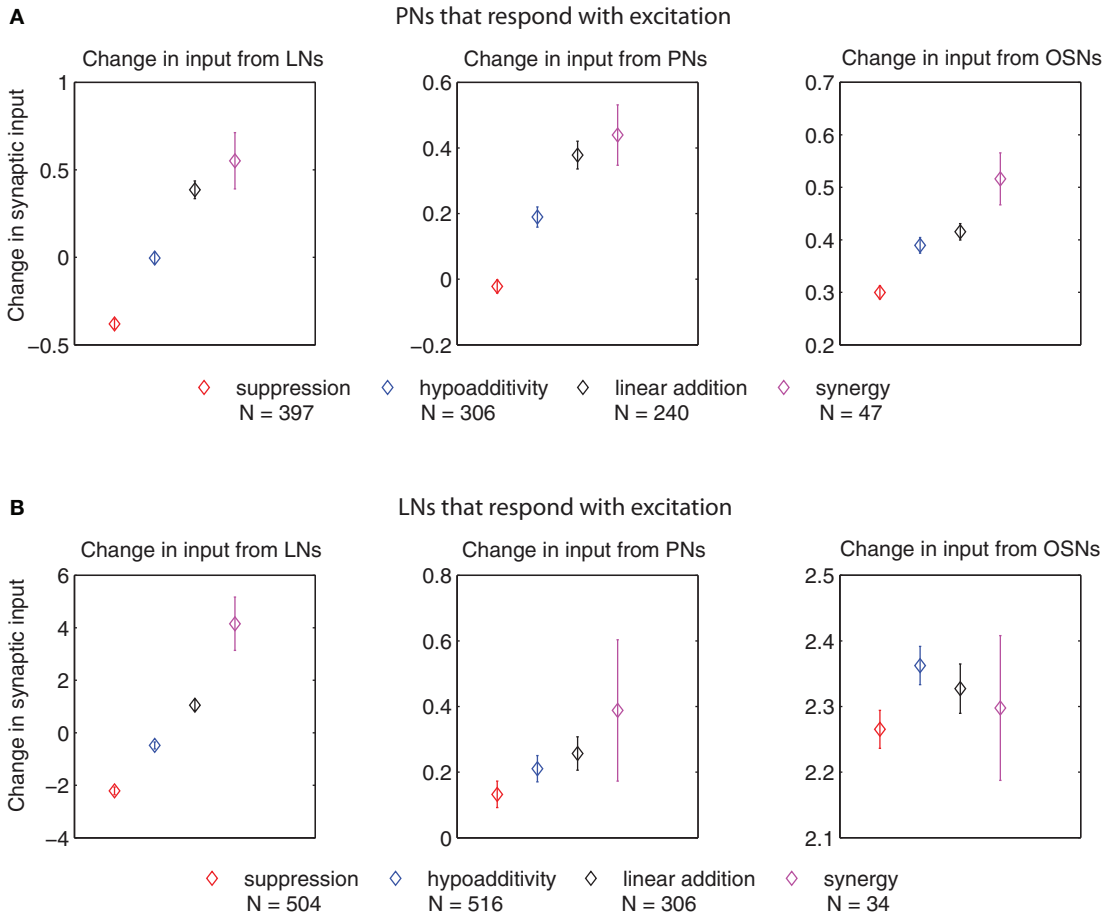

FIGURE 7 | Change in synaptic input (defined as the difference in input between the blend response and the average response to the single components at low concentration, multiplied by the corresponding synaptic weights) for PNs and LNs that respond with excitation of different blend interaction types ($\mu \pm$ SEM for 100 model realizations). **(A)** Change in input to PNs arriving from LNs (left panel), PNs (middle panel), and OSNs (right panel). **(B)** Change in input to LNs arriving from LNs (left panel), PNs (middle panel), and OSNs (right panel). The p values for cases in which we found significant differences in the t-test ($p \leq 0.01$, after Bonferroni correction) are listed in the following. For the change in input arriving from LNs to PNs that respond with excitation (left panel in **A**): suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p < 1 \times 10^{-6}$, suppression vs. synergy $p < 1 \times 10^{-6}$; hypoadditivity vs. linear addition $p < 1 \times 10^{-6}$; hypoadditivity vs. synergy $p = 1 \times 10^{-6}$. For the change in input arriving from PNs to PNs that respond with excitation (middle panel in **A**): suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p < 1 \times 10^{-6}$; suppression vs. synergy $p < 1 \times 10^{-6}$; hypoadditivity vs. linear addition $p < 1.41 \times 10^{-3}$.For the change in input arriving from OSNs to PNs that respond with excitation (right panel in **A**) suppression vs. hypoadditivity $p = 2.8 \times 10^{-5}$; suppression vs. linear addition $p < 1 \times 10^{-6}$; suppression vs. synergy $p = 1 \times 10^{-6}$. For the change in input arriving from LNs to LNs that respond with excitation (left panel in **B**): suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p < 1 \times 10^{-6}$; suppression vs. synergy $p < 1 \times 10^{-6}$; hypoadditivity vs. linear addition $p =< 1 \times 10^{-6}$, hypoadditivity vs. synergy $p < 1 \times 10^{-6}$, linear addition vs. synergy $p = 7.9 \times 10^{-5}$.

neuroscience and is in agreement with most observations of the activation profile in real neurons.

Our computational model of the AL comprises many levels of biological organization, from the OSN periphery to the network architecture, including the neuronal model and synaptic interactions. We used probabilistic connectivity to create random networks that operate in a balanced regime, with strong excitation and strong inhibition that approximately compensate each other, which is a reasonable strategy considering that the details of the local synaptic efficacy that operates in the biological system are unknown. This allows to reproduce the proportion of the different interaction type even when some key parameters (such as the lateral excitation and lateral inhibition synaptic weights; see "Network connectivity of the model") are varied in a relatively broad range. Whenever possible, morphological and physiological data were used to constrain the model, but in some cases simplifying assumptions had to be made due to the lack of sufficient experimental evidence. In such cases we adopted decisions that were plausible or based on previous literature. For the periphery, a linear combination of ingredients was used because it is the simplest assumption and no information is available about this issue in the moth. The neuronal and synaptic models were taken from a previous modeling study (Chong et al., 2012), and the probabilistic rules used to set the network structure were

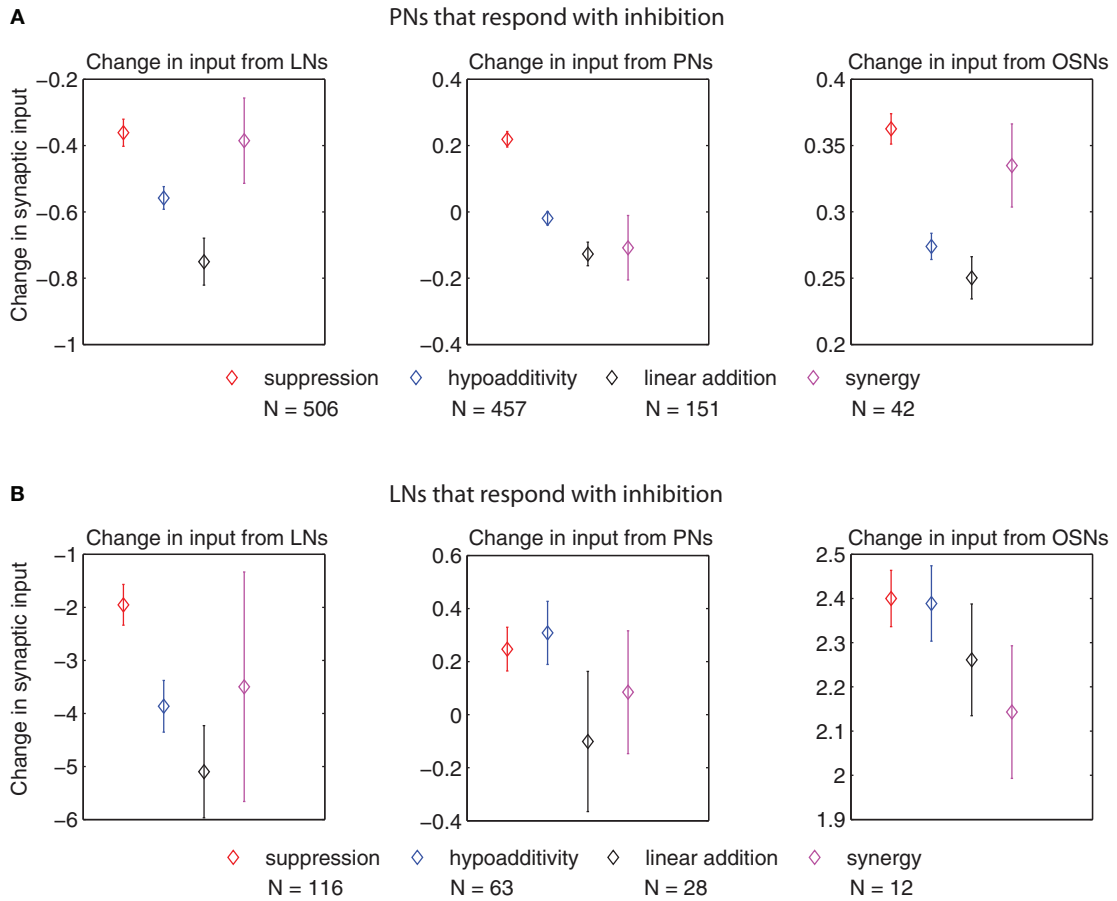

FIGURE 8 | Change in synaptic input (defined as the difference in input between the blend response and the average response to the single components at low concentration, multiplied by the corresponding synaptic weights) for PNs and LNs that respond with inhibition of different blend interaction types (μ ± SEM for 100 model realizations). **(A)** Change in input to PNs arriving from LNs (left panel), PNs (middle panel), and OSNs (right panel). **(B)** Change in input to LNs arriving from LNs (left panel), PNs (middle panel), and OSNs (right panel). The p values for cases in which we found significant differences in the t-test ($p \leq 0.01$, after Bonferroni correction) are listed in the following. For the change in input arriving from LNs to PNs that respond with inhibition (left panel in **A**): suppression vs. hypoadditivity $p = 1.64 \times 10^{-3}$; suppression vs. linear addition $p = 2.9 \times 10^{-5}$. For the change in input arriving from PNs to PNs that respond with inhibition (middle panel in **A**): suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p < 1 \times 10^{-6}$; suppression vs. synergy $p = 9.78 \times 10^{-4}$. For the change in input arriving from OSNs to PNs that respond with inhibition (right panel in **A**): suppression vs. hypoadditivity $p < 1 \times 10^{-6}$; suppression vs. linear addition $p = 6 \times 10^{-6}$. For the change in input arriving from LNs to LNs that respond with inhibition (left panel in **B**): suppression vs. linear addition $p = 3.3 \times 10^{-3}$.

based on morphological studies (Kuebler et al., 2011, and references therein). The parameters are biologically plausible and allow us to reproduce the excitation/inhibition ratio and response sparsity observed in the recordings. Our results are based on the average behavior of 100 networks generated with probabilistic connectivity rules, hence the possibility that they rely on a specific value for one or more of the non-ranged, fixed parameters is extremely unlikely. The general agreement between the proportions of blend interactions in the simulations and the recordings was not calibrated into the model but arises as a purely emergent phenomenon. This predictive power strongly indicates that the model describes the mean responses of AL neurons to both pure chemicals and odor blends in a physiologically relevant manner, capturing the working principles of the AL network to a considerable extent.

In the following, we discuss the role of synaptic input within the network in determining the type of response and blend interaction displayed by individual neurons of the model. Our analysis is restricted to the mean values of neural activity during the odor stimulus and control period, leaving the influence of the dynamic patterns to be presented in a separate study.

SYNAPTIC INPUT, RESPONSE TYPES, AND BLEND INTERACTIONS IN THE COMPUTATIONAL MODEL

Neurons that respond with excitation undergo an increase in net synaptic input in response to blend stimulation, while neurons that respond with inhibition experience a decrease in their input, as expected (**Figure 5A**). Individual neurons are implemented as leaky integrators, hence they can only increase (decrease) their activation if their input increases (decreases). A more

surprising result is that the difference in total input depicted in **Figure 5A** is determined mainly by the input arriving from the LNs (**Figures 5B,C**). The LNs that respond with inhibition do so because they received a very strong inhibition from other LNs, and not less excitation from the PNs and OSNs. The LNs that respond with excitation do so because they received weaker inhibition from the LNs and not larger excitation from the PNs and OSNs (**Figure 5C**). In the case of the PNs, the response consisting of inhibition was built up by a very strong inhibition from LNs combined with a weaker excitation from PNs and OSNs, while for the response consisting of excitation a weaker inhibition from LNs was combined with a stronger excitation from PNs and OSNs (**Figure 5B**). To summarize this first finding, we can state that the response type of all neurons depends on the change in LN input in response to olfactory stimulation, with an agonistic contribution of smaller magnitude coming from the PN and OSN input in the case of PNs.

The second finding in our assessment of the pre-synaptic activity (**Figures 6**, **7**, and **8**), is that the change in synaptic input from the LNs is key to determining the type of blend interaction shown by a neuron (PN or LN). In the particular case of LNs that respond with excitation, the change in input from other LNs is the only determinant of the blend interaction type, with suppression receiving the largest amount of inhibition (i.e., lower input value) followed by hypoadditivity, linear addition, and synergy in a sequence of decreasing inhibition. A similar situation holds for PNs that respond with excitation, where this sequence starts in suppression, continues with hypoadditivity, and ends with linear addition (without including synergy). However, the change in input arriving from other PNs also contributes to determine the blend interaction type through an excitation drive that progressively grows in the sequence suppression, hypoadditivity, linear addition, and synergy. In addition, the change in input from OSNs to PNs is smaller in suppression than in the other categories, creating an additional influence that adds agonistically with the input from the LNs and PNs. In the case of neurons that respond with inhibition we can also recognize a key role of the LNs, although in this case it is mainly to determine the difference between suppression and other categories, without the sequence mentioned above. Again, for LNs that respond with inhibition the only significant influence is the one coming from other LNs, while for PNs that respond with inhibition the change in input arriving from PNs and OSNs combines in an agonistic manner with the change in LN input to determine whether the interaction type will be suppression or not. From the three agonistic influences observed in PNs that respond with excitation and inhibition, the one coming from the LNs is stronger as it can be seen in the scales of ordinates in the **Figures 7** and **8**.

As a summary of our synaptic analysis, we can state that the input changes coming from LNs, PNs, and OSNs drive the membrane potential of the PN output from the AL in an agonistic manner that determines the type of blend interaction that they display. In the case of the LNs, the only significant input is the one arriving from other LNs. Suppressive interactions in neurons that respond with excitation are associated with smaller changes in synaptic input than the other types of blend interaction, while the opposite holds for neurons that respond with inhibition.

The observed proportions of blend interactions types were robust to changes in lateral excitation, but were affected by decreases in the weight of lateral inhibition. The fact that less LN to LN coupling decreased the proportion of suppression (increasing linear addition by a similar amount; see "High correlation in blend interaction types between physiological data and computational model") in neurons that respond with excitation makes sense on the light of the analysis presented in **Figure 7B** (left panel), as suppressive interactions are associated with negative values of synaptic input change from LNs. Thus, lateral inhibition appears to be playing a more important role than the lateral excitation (mediated by the PN to PN connectivity) in the shaping of blend interactions types. This agrees with the results of the honey bee modeling studies of Linster and Smith (1997) for suppression and hypoadditivity (referred in Linster's study as blocking and overshadowing, respectively) and Schmuker et al. (2011) who showed that strong lateral inhibition can result in suppressive mixture coding, allowing good odor discrimination in the PNs of the lateral antenno-cerebral tract. Our results suggest that lateral inhibition is also very important for the emergence of linear addition and synergism in the case of the neurons that respond with excitation (left panels in **Figures 7A,B**), since these interaction types are associated with a larger input change arriving from LNs. Behavioral studies have shown that $GABA_A$ antagonists disrupt odor discrimination (Mwilaria et al., 2008) and the bursting response pattern of PNs associated with odor source location (Lei et al., 2009) in *M. sexta*. This indicates that pharmacological or genetic manipulation of the interneuron network of the moth would result in a severe reduction in non-linear interactions to blends, as has been observed for mixture suppression in the PNs of *Drosophila* after picrotoxin application (Silbering and Galizia, 2007). This type of experiments, combined with simultaneous multi-unit and optical recordings that assess a greater proportion of the AL network, would be key to elucidating the nature of AL blend processing mechanisms in the future.

RELATION WITH PREVIOUS EXPERIMENTAL AND MODELING STUDIES

The results of this computational study provide important global insights into the biological network that are difficult to uncover empirically through current physiological methods. Comparative analyses across several species suggest that odor blends are coded at the first processing stage in a spatiotemporal fashion (for review see Lei and Vickers, 2008) defined by a stereotyped spatial pattern (Galizia and Roessler, 2010), different response onsets (Krofczik et al., 2008), and synchronous ensemble firing patterns (Riffell et al., 2009a,b). However, the close agreement between our neuromorphic model and the electrophysiological data suggests that the connectivity pattern of input OSNs, interglomerular LNs, and output PNs is itself sufficient to confer the level of blend interactions exhibited by the AL. This implies that other cellular characteristics relating to the firing and spatiotemporal dynamics of AL neurons are not obligatory to establish the non-linearity of blend processing witnessed in other electro- and optophysiological analyses of insects (Carlsson et al., 2005; Deisig et al., 2006; Pinero et al., 2008; Silbering et al., 2008; Riffell et al., 2009a). This result is particularly surprising considering the highly reduced assumptions made in our model that

do not consider cellular and network properties such as excitatory interneuron synapses (Shang et al., 2007), heterogeneity of interneuron branching (Fonta et al., 1993), and additional electric coupling between glomeruli (Yaksi and Wilson, 2010).

Further examination of the synaptic interactions in the model reveals that the source of the non-linear processing is due in large part to the interneurons within the network. Such a result is reasonable when one considers the broadly tuned nature of the input to the AL. A blend response results from the activity of several OSNs with different affinities and responses to odor components (Hallem et al., 2004; Hallem and Carlson, 2006). As receptor neurons expressing the same receptor innervate the same glomerulus in the AL (Ressler et al., 1994; Mombaerts, 1996; Vosshall, 2000), some form of interglomerular connectivity is imperative for accurate blend representation, and any apparent non-linearity is subsequently the result of such interconnectivity. Indeed, recent physiological studies in flies (Silbering et al., 2008) and bees (Deisig et al., 2006) have highlighted the vital role of the interneuron network in conferring blend non-linearity. Moreover, the importance of inhibitory sub-networks in shaping the activity of excitatory neurons and synchronous firing is widely recognized in the field of sensory perception [e.g., see Assisi et al., 2011; but note the recent study of Daly et al. (2011) in *M. sexta*], and our study suggests it is also vital for neural encoding of complex odor blends.

Interestingly, the output of the network, relayed by PNs, is ultimately influenced by synaptic input from all three types of AL neurons. This indicates that the final representation of a blend is a composite of all possible interactions within the AL. This also indicates that any blend non-linearities already present at the periphery (not included here but implied by other studies: Carlsson and Hansson, 2002; Hillier and Vickers, 2011; Su et al., 2011) may significantly impact the resultant output from the AL. Additional modulation between PNs (either within or between glomeruli) further modifies the output and creates the ultimate "blend percept" (Kuebler et al., 2011).

Our results suggest that the non-linear processing establishing the unique "blend percept" within the AL can result from network interactions, without the need of intrinsic neuronal properties in the cells within that network. Our findings highlight that such mixture interactions are a natural outcome of the architecture of the AL, and reveal its important role in shaping the perception of olfactory information in the CNS. By design, the AL is not merely a relay station for olfactory information, but filters and processes multicomponent information into a unique representation that reduces signal dimensionality for subsequent processing in the CNS.

ACKNOWLEDGMENTS

This research was supported by EU 6th Framework Programme FET Project iCHEM and the Max Planck Society. The authors want to thank Dr. Manuel Montanes for his useful comments and suggestions.

REFERENCES

Assisi, C., Stopfer, M., and Bazhenov, M. (2011). Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning. *Neuron* 69, 373–386.

Baker, T. C. (2008). Balanced olfactory antagonism as a concept for understanding evolutionary shifts in moth sex pheromone blends. *J. Chem. Ecol.* 34, 971–981.

Bruyne, M., and Baker, T. C. (2008). Odor detection in insects: volatile codes. *J. Chem. Ecol.* 34, 882–897.

Carlsson, M. A., Chong, K. Y., Daniels, W., Hansson, B. S., and Pearce, T. C. (2007). Component information is preserved in glomerular responses to binary odor mixtures in the moth *Spodoptera littoralis*. *Chem. Senses* 32, 433–443.

Carlsson, M. A., and Hansson, B. S. (2002). Responses in highly selective sensory neurons to blends of pheromone components in the moth *Agrotis segetum*. *J. Insect Physiol.* 48, 443–451.

Carlsson, M. A., and Hansson, B. S. (2003). Dose-response characteristics of glomerular activity in the moth antennal lobe. *Chem. Senses* 28, 269–278.

Carlsson, M. A., Knusel, P., Verschure, P. F. M. J., and Hansson, B. S. (2005). Spatio-temporal Ca2+ dynamics of moth olfactory projection neurones. *Eur. J. Neurosci.* 22, 647–657.

Chong, K. Y., Capurro, A., Karout, S., and Pearce, T. C. (2012). Stimulus and network dynamics can collide in a ratiometric model of the antennal lobe macroglomerular complex. *PLoS One* 7:e29602. doi: 10.1371/journal.pone.0029602

Christensen, T. A., Mustaparta, H., and Hildebrand, J. G. (1991). Chemical communication in heliothine moths. II. Central processing of intraspecific and interspecific olfactory messages in the male cOSN earworm moth, *Helicoverpazea*. *J. Comp. Physiol. A* 169, 259–274.

Daly, K. C., Galán, R. F., Peters, O. J., and Staudacher, E. M. (2011). Detailed characterization of local field potential oscillations and their relationship to spike timing in the antennal lobe of the moth *Manduca sexta*. *Front. Neuroeng.* 4:12. doi: 10.3389/fneng.2011.00012

Deisig, N., Giurfa, M., Lachnit, H., and Sandoz, J. C. (2006). Neural representation of olfactory mixtures in the honeybee antennal lobe. *Eur. J. Neurosci.* 24, 1161–1174.

Deisig, N., Giurfa, M., and Sandoz, J. C. (2010). Antennal lobe processing increases separability of odor mixture representations in the honeybee. *J. Neurophysiol.* 103, 2185–2194.

Duchamp-Viret, P., Duchamp, A., and Chaput, M. (2003). Single olfactory sensory neurons simultaneously integrate the components of an odour mixture. *Eur. J. Neurosci.* 18, 2690–2696.

Fernandez, P. C., Locatelli, F. F., Person-Rennell, N., Deleo, G., and Smith, B. H. (2009). Associative conditioning tunes transient dynamics of early olfactory processing. *J. Neurosci.* 29, 10191–10202.

Fonta, C., Sun, X., and Masson, C. (1993). Morphology and spatial distribution of bee antennal lobe interneurones responsive to odours. *Chem. Senses* 18, 101–119.

Galizia, C., and Menzel, R. (2001). The role of glomeruli in the neural representation of odours: results from optical recording studies. *J. Insect Physiol.* 47, 115–130.

Galizia, C. G., and Roessler, W. (2010). Parallel olfactory systems in insects: anatomy and function. *Annu. Rev. Entomol.* 55, 399–420.

Grosse-Wilde, E., Kuebler, L. S., Bucks, S., Vogelb, H., Wichera, D., and Hansson, B. S. (2011). Antennal transcriptome of *Manduca sexta*. *Proc. Natl. Acad. Sci. U.S.A.* 108, 7449–7454.

Hallem, E. A., and Carlson, J. R. (2006). Coding of odors by a receptor repertoire. *Cell* 125, 143–160.

Hallem, E. A., Ho, M. G., and Carlson, J. R. (2004). The molecular basis of odor coding in the *Drosophila antenna*. *Cell* 117, 965–979.

Hansson, B. S., and Anton, S. (2000). Function and morphology of the antennal lobe: new developments. *Annu. Rev. Entomol.* 45, 203–231.

Hillier, N. K., and Vickers, N. J. (2011). Mixture interactions in moth olfactory physiology: examining the effects of odorant mixture, concentration, distal stimulation, and antennal nerve transection on sensillar responses. *Chem. Senses* 36, 93–108.

Homberg, U., Christensen, T., and Hildebrand, J. (1989). Structure and function of the deutocerebrum in insects. *Annu. Rev. Entomol.* 34, 477–501.

Homberg, U., Montague, R. A., and Hildebrand, J. G. (1988). Anatomy of antenno-cerebral pathways in the brain of the sphinx moth *Manduca sexta*. *Cell Tissue Res.* 254, 255–281.

Huang, J., Zhang, W., Qiao, W., Hu, A., and Wang, Z. (2010).

Functional connectivity and selective odor responses of excitatory local interneurons in Drosophila antennal lobe. *Neuron* 67, 1021–1033.

Joerges, J., Küttner, A., Galizia, C., and Menzel, R. (1997). Representations of odours and odour mixtures visualized in the honeybee brain. *Nature* 387, 285–288.

Krofczik, S., Menzel, R., and Nawrot, M. P. (2008). Rapid odor processing in the honeybee antennal lobe network. *Front. Comput. Neurosci.* 2:9. doi: 10.3389/neuro.10.009.2008

Kuebler, L. S., Olsson, S. B., Weniger, R., and Hansson, B. S. (2011). Neuronal processing of complex mixtures establishes a unique odor representation in the moth antennal lobe. *Front. Neural Circuits* 5:7. doi: 10.3389/fncir.2011.00007

Kuebler, L. S., Schubert, M., Karpati, Z., Hansson, B. S., and Olsson, S. B. (2012). Antennal lobe processing correlates to moth olfactory behavior. *J. Neurosci.* 32, (in press).

Laing, D. G., Panhuber, H., and Slotnik, B. M. (1989). Odour masking in the rat. *Physiol. Behav.* 45, 689–694.

Lei, H., Riffel, J. A., Gage, S. L., and Hildebrand, J. G. (2009) Contrast enhancement of stimulus intermittency in a primary olfactory network and its behavioral significance. *J. Biol.* 8, 21.1–21.16.

Lei, H., and Vickers, N. (2008). Central processing of natural odor mixtures in insects. *J. Chem. Ecol.* 34, 915–927.

Linster, C., and Cleland, T. A. (2010). Decorrelation of odor representations via spike timing-dependent plasticity. *Front. Comput. Neurosci.* 4:157. doi: 10.3389/fncom.2010.00157

Linster, C., Sachse, S., and Galizia, C. G. (2005). Computational modeling suggests that response properties rather than spatial position determine connectivity between olfactory glomeruli. *J. Neurophysiol.* 93, 3410–3417.

Linster, C., and Smith, B. H. (1997). A computational model of the response of honey bee antennal lobe circuitry to odor mixtures: overshadowing, blocking and unblocking can arise from lateral inhibition. *Behav. Brain Res.* 87, 1–14.

Martin, J. P., Beyerlein, A., Dacks, A. M., Reisenman, C. E., Riffell, J. A., Lei, H., and Hildebrand, J. G. (2011). The neurobiology of insect olfaction: sensory processing in a comparative context. *Prog. Neurobiol.* 95, 427–447.

Matsumoto, S., and Hildebrand, J. (1981). Olfactory mechanisms in the moth *Manduca sexta*: response characteristics and morphology of central neurons in the antennal lobes. *Proc. R. Soc. Lond. B Biol. Sci.* 213, 249–277.

Meyer, A., and Galizia, C. G. (2011) Elemental and con?gural olfactory-coding by antennal lobeneurons of the honey bee (*Apis mellifera*). *J. Comp. Physiol. A* 198, 159–171.

Mombaerts, P. (1996). Targeting olfaction. *Curr. Opin. Neurobiol.* 6, 481–486.

Mwilaria, E. K., Ghatak, C., and Daly, K. C. (2008). Disruption of GABAA in the insect antennal lobe generally increases odor detection and discrimination thresholds. *Chem. Senses* 33, 267–281.

Olsson, S., Kuebler, L., Veit, D., Steck, K., Schmidt, A., Knaden, M., and Hansson, B. (2011). A novel multicomponent stimulus device for use in olfactory experiments. *J. Neurosci. Methods* 195, 1–9.

Pinero, J. C., Galizia, C. G., and Dorn, S. (2008). Synergistic behavioral responses of female oriental fruit moths (Lepidoptera:Tortricidae) to synthetic host plant-derived mixtures are mirrored by odor-evoked calcium activity in their antennal lobes. *J. Insect Physiol.* 54, 333–343.

Rabinovich, M., Huerta, R., and Laurent, G. (2008). Transient dynamics for neural processing. *Science* 321, 48–50.

Reisenman, C. E., Dacks, A. M., and Hildebrand, J. G. (2011) Local interneuron diversity in the primary olfactory center of the moth *Manduca sexta*. *J. Comp. Physiol. A* 197, 653–665.

Ressler, K. J., Sullivan, S. L., and Buck, L. B. (1994). A molecular dissection of spatial patterning in the olfactory system. *Curr. Opin. Neurobiol.* 4, 588–596.

Riffell, J. A., Lei, H., Christensen, T. A., and Hildebrand, J. G. (2009a). Characterization and coding of behaviorally significant odor mixtures. *Curr. Biol.* 19, 335–340.

Riffell, J. A., Lei, H., and Hildebrand, J. G. (2009b). Neural correlates of behavior in the moth *Manduca sexta* in response to complex odors. *Proc. Natl. Acad. Sci. U.S.A.* 106, 19219–19226.

Rospars, J. P., Lansky, P., Chaput, M., and Duchamp-Viret, P. (2008). Competitive and noncompetitive odorant interactions in the early neural coding of odorant mixtures. *J. Neurosci.* 28, 2659–2666.

Schmuker, M., Yamagata, N., Nawrot, M., and Menzel, R. (2011). Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee. *Front. Neuroeng.* 4:17. doi: 10.3389/fneng.2011.00017

Shang, Y., Claridge-Chang, A., Sjulson, L., Pypaert, M., and Miesenböck, G. (2007). Excitatory local circuits and their implications for olfactory processing in the fly antennal lobe. *Cell* 128, 601–612.

Silbering, A. F., and Galizia, C. G. (2007). Processing of odor mixtures in the Drosophila antennal lobe reveals both global inhibition and glomerulus-specific interactions. *J. Neurosci.* 27, 11966–11977.

Silbering, A. F., Okada, R., Ito, K., and Galizia, G. (2008). Olfactory information processing in the Drosophila antennal lobe: anything goes? *J. Neurosci.* 28, 13075–13087.

Su, C.-Y., Martelli, C., Emonet, T., and Carlson, J. R. (2011). Temporal coding of odor mixtures in an olfactory receptor neuron. *PNAS*, 108, 5075–5080.

Vosshall, L. (2000). Olfaction in Drosophila. *Curr. Opin. Neurobiol.* 10, 498–503.

Yaksi, E., and Wilson, R. I. (2010). Electrical coupling between olfactory glomeruli. *Neuron* 67, 1034–1047.

Yamagata, N., Schmuker, M., Szyszka, P., Mizunami, M., and Menzel, R. (2009). Differential odor processing in two olfactory pathways in the honeybee. *Front. Syst. Neurosci.* 3:16. doi: 10.3389/neuro.06.016.2009

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.

Received: 02 November 2011; accepted: 14 March 2012; published online: 19 April 2012.

*Citation: Capurro A, Baroni F, Olsson SB, Kuebler LS, Karout S, Hansson BS and Pearce TC (2012) Non-linear blend coding in the moth antennal lobe emerges from random glomerular networks. Front. Neuroeng. **5**:6. doi: 10.3389/fneng.2012.00006*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 03 January 2013
doi: 10.3389/fneng.2012.00019

Network architecture underlying maximal separation of neuronal representations

Ron A. Jortner*

Interdisciplinary Center for Neural Computation, Hebrew University, Jerusalem, Israel

Edited by:
Thomas Nowotny, University of Sussex, UK

Reviewed by:
Pentti Kanerva, NASA Ames Research Center, USA (Retired)
Thomas Nowotny, University of Sussex, UK
Amir Madany, University of Luebeck, Germany

****Correspondence:***
Ron A. Jortner, Department of Cellular and Systems Neurobiology, Max Planck Institute of Neurobiology, Am Klopferspitz 18, 82152 Martinsried, Germany.
e-mail: ronijort@neuro.mpg.de

One of the most basic and general tasks faced by all nervous systems is extracting relevant information from the organism's surrounding world. While physical signals available to sensory systems are often continuous, variable, overlapping, and noisy, high-level neuronal representations used for decision-making tend to be discrete, specific, invariant, and highly separable. This study addresses the question of how neuronal specificity is generated. Inspired by experimental findings on network architecture in the olfactory system of the locust, I construct a highly simplified theoretical framework which allows for analytic solution of its key properties. For generalized feed-forward systems, I show that an intermediate range of connectivity values between source- and target-populations leads to a combinatorial explosion of wiring possibilities, resulting in input spaces which are, by their very nature, exquisitely sparsely populated. In particular, connection probability ½, as found in the locust antennal-lobe–mushroom-body circuit, serves to maximize separation of neuronal representations across the target Kenyon cells (KCs), and explains their specific and reliable responses. This analysis yields a function expressing response specificity in terms of lower network parameters; together with appropriate gain control this leads to a simple neuronal algorithm for generating arbitrarily sparse and selective codes and linking network architecture and neural coding. I suggest a straightforward way to construct ecologically meaningful representations from this code.

Keywords: neural coding, sparseness, circuit, connectivity, specificity, olfaction, insect, locust

INTRODUCTION

Animals all use information about their surrounding world in order to function within it. Nervous systems have specialized in gathering, processing, storing, and retrieving such information and in using it to make decisions necessary for survival. To accomplish these tasks, the brain must disregard much of the information made available by the senses, extracting only what is relevant for the animal's needs. Just as in drawing a map of a newly discovered land, the brain, in so doing, creates a schematic internal representation of the animal's world—and it is over this internal model that generalizations are drawn, categories are discerned, associations made, and behavior triggered [Marr, 1970, 1971; Barlow, 1985; von der Malsburg, 1986, 1990; Kanerva, 1988; Földiák, 1990; reviewed in deCharms and Zador (2000)].

By virtue of the choice of what to keep in it, this internal neuronal representation is tailored to the organism's needs; and just as a historian, geologist, and meteorologist would each draw a different map of the same piece of land, it too suggests alternate ways of viewing and interpreting reality (Barlow, 1972; Kanerva, 1988; Churchland and Sejnowski, 1990). In other words, a subjective internal model of the world serves as a substrate for performing computations which—by predicting the outcome of actions in the real world—allow efficient decision-making, even in novel situations (von der Malsburg, 1986, 1990; Kanerva, 1988). This may be the core of what the brain does.

Olfactory systems, which are in evolutionary terms ancient and found even in simple animals, accomplish this task very efficiently. The signals they analyze are plumes of airborne molecules and complex mixtures thereof—variable signals occurring on highly noisy background (Kadohisa and Wilson, 2006; Raman and Stopfer, 2010; Raman et al., 2011)—and from this input they extract meaning (such as "food," "predator," or "potential sexual partner"), which is translated into behavioral output (actions such as foraging, escape, or courtship, respectively).

How is this task accomplished by neural hardware? Circuit architecture is a key to understand brain dynamics and function. A full characterization of neural circuitry—including cell types and their integrative properties (input–output functions), connectivity between them (statistics, pattern, signs, and strengths) and external input driving the network (rates, auto- and cross-correlations, synchrony, etc.)—is necessary, though not sufficient, for transcending the descriptive level and distilling the system's design principles (Churchland and Sejnowski, 1992). This in turn yields a deeper understanding of how basic network features and their interrelations give rise to its higher properties. Few biological neural systems, however, are presently characterized in sufficient detail; most are riddled with complexity, knowledge gaps, and high-dimensional parameter-spaces.

One example where detailed knowledge exists on network parameters and coding schemes is the olfactory system of the

locust (*Schistocerca americana*) (**Figure 1A**). In this relatively simple system, 800 broadly tuned and noisy second-order neurons (projection neurons, PNs) project directly onto 50,000 third-order neurons (Kenyon cells, KCs), which are highly selective and reliable in their odor responses (Perez-Orive et al., 2002). As the system is feed-forward, small, well-defined, and displays a dramatic change in coding—from distributed to sparse—between source- and target-populations, it seems well suited for studying the origins of neuronal specificity.

The locust olfactory system (**Figure 1A**) receives odor input through the antenna, via ∼90,000 olfactory receptor-neurons (ORNs) which terminate in the antennal-lobe. The antennal-lobe is a small network: ∼800 excitatory PNs which send their axons to the next relays in the system (forming the antennal-lobe's sole output), and ∼300 inhibitory interneurons (not shown in the diagram) which act locally within the network (Laurent and Davidowitz, 1994; Leitch and Laurent, 1996). PNs each respond to a wide array of odors with rich, complex spike-trains encoding odor identity (Laurent and Davidowitz, 1994; Laurent, 1996; Wehr and Laurent, 1996; Perez-Orive et al., 2002; Mazor and Laurent, 2005) and concentration (Stopfer et al., 2003); PN-spike-trains are additionally locked to a 20 Hz oscillatory cycle

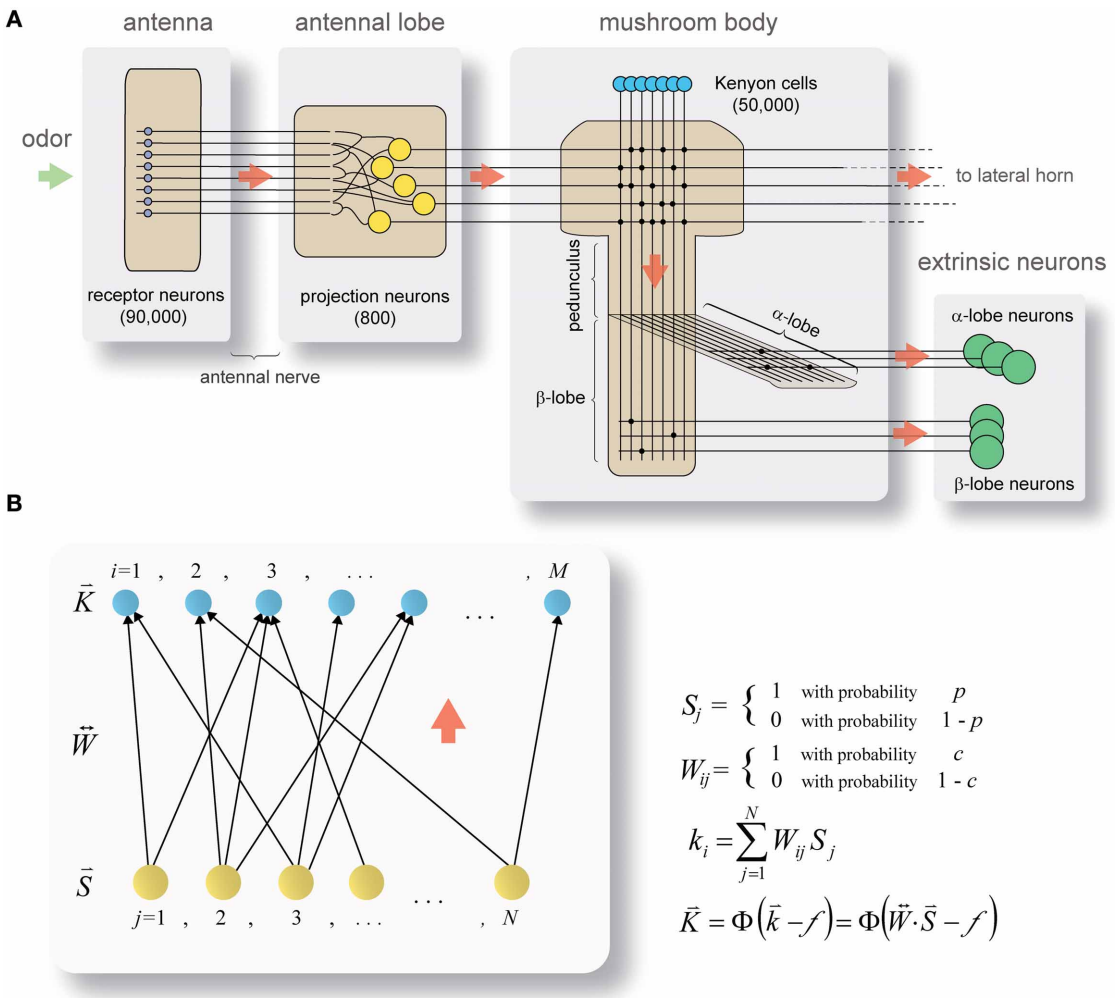

FIGURE 1 | Framework for studying the separation of neuronal representations. (A) Circuit diagram of the locust olfactory system. Odor information reaches the antennal-lobe via ∼90,000 olfactory receptor-neurons in the antenna. In the antennal-lobe, ∼800 projection neurons (PNs, yellow) project it further to the mushroom body (onto ∼50,000 Kenyon cells (KCs), blue; PN–KC connection probability ½) and the lateral horn (not shown). In transition from PNs to KCs the odor code dramatically changes from broad and highly distributed (in PNs) to sparse and specific (in KCs). KC axons split into the α- and β-lobes, where they synapse onto α- and β-lobe extrinsic neurons (green), respectively (KC–β-lobe-neuron connection probability ∼0.02). Red arrows indicate direction of information flow. See text for more details. **(B)** Mathematical framework for studying transformation in coding. Model represents the state of a theoretical network inspired by PN–KC circuitry during a brief snapshot in time. Color code and information flow same as in **(A)**. A set of *N* source-neurons (activity of which is denoted by binary-valued vector $\vec{S}$; i.i.d. with probability *p*) projects onto a set of *M* target neurons (activity of which is denoted by vector $\vec{K}$) via a set of feed-forward connections (binary-valued connectivity matrix $\vec{\vec{W}}$; i.i.d. with probability *c*). The aggregate input to the target layer $\vec{K}$ is the vector $\vec{k}$, a product of source-neuron activity vector $\vec{S}$ and connectivity matrix $\vec{\vec{W}}$. $\vec{K}$ is obtained by thresholding $\vec{k}$ using the Heaviside function $\Phi\left(\vec{k} - f\right)$.

which is synchronous across the PN population (Laurent and Davidowitz, 1994; Laurent, 1996; Laurent et al., 1996) and is reflected in local-field-potential oscillations. With no odor presented, PNs fire spontaneously at rates of 2.5–4 Hz (Perez-Orive et al., 2002; Mazor and Laurent, 2005). Odors are represented by a dynamic combinatorial code (Laurent et al., 1996; Wehr and Laurent, 1996) which is broadly distributed across the PN population (Perez-Orive et al., 2002; Mazor and Laurent, 2005).

Output from the antennal-lobe is projected, via PN axons, onto two direct target-areas: the mushroom body, a structure involved in learning and memory (Heisenberg, 1998), and the lateral horn. The mushroom body contains ∼50,000 small neurons, the KCs (Laurent and Naraghi, 1994; Leitch and Laurent, 1996). Individual KCs respond to specific odors (either monomolecular odors or mixtures), their responses are characterized by few spikes, are highly reliable across different presentations of the same odor (Perez-Orive et al., 2002), and are often concentration invariant (Stopfer et al., 2003). KC responses occur on a background of extremely little spontaneous firing (Laurent and Naraghi, 1994; Perez-Orive et al., 2002; Mazor and Laurent, 2005; Jortner et al., 2007; Jortner, 2009). Mushroom body odor-responses thus involve small, highly selective subsets of KCs (Perez-Orive et al., 2002, 2004; Stopfer et al., 2003; Jortner, 2009).

Axons of KCs exit the mushroom body calyx in a tight bundle (forming the mushroom's stalk, or pedunculus), branching into the mushroom body's output nodes, the α- and β-lobes (Laurent and Naraghi, 1994). There, KC output is integrated by smaller populations of extrinsic neurons (called α- and β-lobe neurons, respectively; **Figure 1A**) with large, planar dendritic trees which intersect KC-axon bundles at neat right angles (Li and Strausfeld, 1997; MacLeod et al., 1998; Cassenaer and Laurent, 2007), suggesting potential integration of precisely timed spikes over a wide KC-subpopulation.

Several previous studies offer theoretical treatment of the locust antennal-lobe–mushroom-body transformation (e.g., Garcia-Sanchez and Huerta, 2003; Theunissen, 2003; Huerta et al., 2004; Sivan and Kopell, 2004; Finelli et al., 2008); these, however, lack quantitative data regarding critical network parameters, such as connectivity values. More recent experimental work quantified aspects of network architecture via electrophysiological measurements of connectivity between PNs, KCs and β-lobe-neurons (Jortner et al., 2007; Cassenaer and Laurent, 2007). Results show that each KC receives synaptic connections from ½ of all PNs on average (∼400 out of ∼800 PNs); PN–KC synapses are very weak [excitatory-postsynaptic-potential (EPSP) amplitude is $85 \pm 44\,\mu V$], and KC firing thresholds correspond to simultaneous activation of ∼100 PN–KC synapses (assuming linear summation) (Jortner et al., 2007). Connections between KCs and some of their outputs (β-lobes neurons) are, on the other hand, sparse (∼2% of pairs) and strong (EPSP amplitude 1.58 ± 1.1 mV), and exhibit Hebbian spike-timing-dependent plasticity (Cassenaer and Laurent, 2007).

Can these findings explain the transformation in coding schemes? What is the functional significance of this design? In the present study I explore design principles by which the brain constructs specific, sparse and high-level representations of the surrounding world. A coding strategy both sparse and selective would be one where *only a small subset of neurons respond to any given stimulus or external state* (i.e., high population sparseness; Willmore and Tolhurst, 2001), *and only a small subset of stimuli or external states elicit response in each neuron* (Jortner et al., 2007; Jortner, 2009). Inspired by the network architecture of the locust olfactory pathways, I suggest an exciting implementation of neuronal hardware to this end. My central claim is that in a feed-forward system with connectivity ½, target neurons differ maximally from each other in information they contain about the world (or external state); in this sense serving as an optimal neural module for parsing the world of inputs, and a substrate for sparse and specific neuronal-responses on the basis of which learning, categorization, generalization, and other essential computations can occur. The targets' sparseness is set to a controlled, arbitrary level by choice of a proper and adaptive firing threshold. Next, I address these points through a straightforward yet rigorous mathematical approach.

METHODS

The model I use is highly reduced, consisting of a layer of source-neurons (equivalent to PNs), projecting onto a layer of target neurons (equivalent to KCs) via a set of feed-forward connections (**Figure 1B**). Following several simplifying assumptions, I describe the mathematical framework and proceed to solve some of its behavior analytically—yielding predictions about function and about how network design relates to coding.

MODEL ASSUMPTIONS

For the sake of tractability and predictive power, I make four important simplifying assumptions. First, I choose to look at a "snapshot" of the system in time; a brief-enough segment so that for any given PN the probability for spiking more than once is negligible. Within this time window, the PN population can be treated as a vector of binary digits, *one* denoting the occurrence of a spike and *zero* denoting none. As a second assumption, all PNs are treated each as firing (or not) within this time window with probability p which is identical across all PNs, and doing so independently of each other (i.i.d.); this allows treating the PN activity vector as binomial with a known parameter. Third, all synaptic connections are treated as equal in strength. As a fourth and last assumption, connectivity between PNs and KCs is assumed to be random, with i.i.d. statistics and probability c across all PN–KC pairs.

These assumptions, and particularly those of i.i.d. statistics of firing and connectivity, wield great predictive power; I will revisit them in the Discussion (Section "Regaining Complexity: Reexamining the Model's Initial Assumptions"), examine their validity with respect to experimental data on the locust olfactory system, and assess, wherever biological reality deviates from them (e.g., when some dependence and correlations are introduced), how model results may be affected.

MODEL DESCRIPTION

A schematic cartoon of the network-model appears in **Figure 1B**. There is a set of N source-neurons, denoted by vector $\vec{S}$ (so the neurons are $S_1, S_2, \ldots, S_N$): these are analogous to PNs in

the antennal-lobe. A second set of M target neurons, denoted by vector $\vec{K}$ (so neurons $K_1, K_2, \ldots, K_M$), are analogous to KCs in the mushroom body. Source-neurons ($\vec{S}$) connect to target neurons ($\vec{K}$) through randomly determined connections of uniform strength; each PN can thus either connect to a given KC or not, with probabilities c and $(1-c)$, respectively. $\overleftrightarrow{W}$ is the connection matrix, with $W_{ij} = 1$ if the jth PN connects to the ith KC and 0 otherwise. Each row of $\overleftrightarrow{W}$ indicates the set of PNs physically connected to a given KC (so there are as many rows as KCs), and each column indicates the set of KCs receiving physical connections from a given PN (so there are as many columns as PNs). The rows I will refer to as the *connectivity vectors* to KCs.

As pointed out in the assumptions, the model looks at a snapshot of the neural system during a brief time window. Within it, each of the PNs can either fire a spike or not, and does so with probabilities p and $(1-p)$, respectively, so $\vec{S}$ also takes binary values. I call $\vec{S}$ the *activity vector* of the PN population, and $\mathbb{S}$ will be the set of all possible activity vectors, so $\vec{S} \in \mathbb{S}$.

Formally, then (**Figure 1B**):

$$S_j = \begin{cases} 1 & \text{with probabil. } p \\ 0 & \text{with probabil. } 1-p \end{cases} \quad j = 1, \ldots, N$$

$$W_{ij} = \begin{cases} 1 & \text{with probabil. } c \\ 0 & \text{with probabil. } 1-c \end{cases} \quad i = 1, \ldots, M;\ j = 1, \ldots, N$$

During our given time window each of the M KCs receives PN inputs, which additively determine its "membrane-potential." The input to each KC, to which I refer throughout this work as its *aggregate input* (denoted by k_i for the ith KC) is the sum of all PNs connected to it which fire during that time window, or formally

$$\vec{k} = \overleftrightarrow{W} \cdot \vec{S}$$
$$k_i = \sum_{j=1}^{N} W_{ij} S_j$$

Thus, $\vec{k}$ is a vector which takes natural values between 0 and N (according to how many of the PNs converging onto the KC fire). Each KC then fires a spike if and only if its aggregate input equals or exceeds the firing threshold, f, or

$$\vec{K} = \Phi\left(\vec{k} - f\right) = \Phi\left(\overleftrightarrow{W} \cdot \vec{S} - f\right)$$
$$K_i = \Phi\left(k_i - f\right) = \Phi\left(\sum_{j=1}^{N} W_{ij} S_j - f\right)$$

where $\Phi(X)$ denotes the Heaviside function:

$$\Phi(X) = \begin{cases} 1 & \text{if } X \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

So $\vec{K}$ is a binary-valued vector, K_i indicating whether or not the ith KC fires, and $\mathbb{K}$ is the set of all possible target-neuron activity vectors, so $\vec{K} \in \mathbb{K}$. Thus, in this model, for a network with N PNs and M KCs (with threshold f) and a fixed connectivity matrix $\overleftrightarrow{W}$, a given state of the PN population (denoted by activity vector $\vec{S}$) unambiguously determines the activity vector of the KC population, $\vec{K}$.

MATHEMATICAL CONVENTIONS, SYMBOLS, AND ABBREVIATIONS

While the mathematics used throughout this work is mostly elementary, some of the derivations are nonetheless rather tedious. For the sake of clarity, they appear in shortened form within the text; I provide commented step-by-step derivations in the Appendix.

All but the most standard mathematical symbols used are defined the first time they appear. For quick reference, they are also listed in **Table 1**.

RESULTS

MODEL RESULTS I: EXPLORING PROPERTIES OF THE CONNECTIVITY MATRIX

Examining the set of connections between the neuronal populations $\vec{S}$ and $\vec{K}$ (connectivity matrix $\overleftrightarrow{W}$), we may ask to what extent two connectivity vectors (rows of $\overleftrightarrow{W}$) differ from each other. Let us calculate how many binary digits will, on average, differ across two such connectivity vectors (which I call $\vec{U}$ and $\vec{V}$). This difference-measure is the *Hamming distance* between the two vectors, denoted by $H(\vec{U}, \vec{V})$. Since all elements of the connectivity matrix are independent from each other, we can simply calculate the probability that an element of $\vec{U}$ differs from the matching element in $\vec{V}$(detailed derivation in Appendix A1):

$$\begin{aligned} \Pr(U_j \neq V_j) &= \Pr(U_j = 1, V_j = 0) + \Pr(U_j = 0, V_j = 1) \\ &= c(1-c) + (1-c)c = 2c(1-c) \end{aligned}$$

and multiply by the total number of elements N to get the Hamming distance:

$$\left\langle H(\vec{U}, \vec{V}) \right\rangle_{\vec{U},\vec{V}} = N \cdot \Pr(U_j \neq V_j) = 2Nc(1-c)$$

As this expression shows, when viewed as a function of the connection probability, c, the Hamming distance between two rows of $\overleftrightarrow{W}$ is maximal for $c = ½$, and drops symmetrically around it (**Figure 2A**, for $N = 800$). Thus, under the model assumptions, PN–KC connectivity vectors will on average be maximally different (as measured by Hamming distance) from each other when each pair of cells (PN and KC) is equally likely to be connected or not. This already suggests some special property of the experimentally observed connectivity matrix (Jortner et al., 2007).

If we now pick two connectivity vectors at random, what is the probability that they are identical? In other words, what is the

Table 1 | Mathematical symbols used throughout the paper.

c	Probability of PN–KC connection (scalar, real within interval [0,1])
CDF	Cumulative Distribution Function
CLT	Central Limit Theorem
$D(x, y)$	Absolute difference between x and y, $\|x - y\|$ (For binary values: $(x - y)^2$)
f	Kenyon-cell firing threshold (in units of PN inputs)—(scalar, non-negative)
$H(\vec{X}, \vec{Y})$	Hamming distance between binary-valued vectors $\vec{X}$, $\vec{Y}$; number of bits by which they differ (scalar, real, non-negative)
i.i.d.	independent and identically distributed
$\vec{K}$	Activity vector of the Kenyon-cell population (vector, $Mx1$, binary values)
$\vec{k}$	vector of aggregate inputs to Kenyon cells (vector, $Mx1$, natural values)
$\mathbb{K}$	Set of all possible Kenyon-cell activity vectors
M	Total number of Kenyon-cells (scalar, natural)
N	total number of PNs (scalar, natural)
p	PN-firing probability within characteristic time window (scalar, real within interval [0,1])
PDF	Probability Density Function
$\Pr(x)$	Probability of x
Q; $Q(x)$	The Standard Normal cumulative distribution function; its value at x
$\vec{S}$	Activity vector of the PN population (vector, $Nx1$, binary values)
$\mathbb{S}$	Set of all possible PN activity vectors
$\vec{U}, \vec{V}$	Random PN–KC connectivity vectors; rows of $\overleftrightarrow{W}$ (vectors, 1 x N, binary values)
u,v	Random subsets of PNs
$\overleftrightarrow{W}$	Connectivity matrix between PNs and KCs (Matrix, M x N, Binary values)
Λ	variance of aggregate input to a KC (scalar, non-negative)
$\Phi(x)$	The Heaviside (step) function; producing 1 if $x \geq 0$ and 0 otherwise
Ψ	Mean aggregate input to a KC (scalar, real)
$\rho(x, y)$	Pearson's correlation coefficient between x and y
$\sim$	Equals in distribution
$\equiv$	Equals by definition
$A \bigcap B$	Intersection of sets A and B (objects which belong to both A and B)
$A \bigcup B$	Union of sets A and B (objects which belong to A or B, *inclusive or*)
$A \Delta B$	Symmetric difference of sets A and B (objects belong to A or B, *but not both*)
$\|\|A\|\|$	Number of elements in set A
$A \backslash B$	Relative complement of sets A and B (objects belong to A and not to B)
$x!$	Factorial of x
$\|x\|$	Absolute value of x
$\langle X \rangle_Y$	Expected value of X over all possible values of Y (with their respective probabilities)
$\begin{bmatrix} x \\ y \end{bmatrix}$	x-choose-y, the number of ways to pick y elements out of x

probability that two randomly chosen KCs sample the exact same ensemble of PNs?

$$\begin{aligned}\Pr\left(H\left(\vec{U},\vec{V}\right)=0\right) &= \Pr\left(\sum_{j=1}^{N}(U_j - V_j)^2 = 0\right)\\ &= \Pr\left(U_j = V_j | \forall j\right)\\ &= \left(\Pr\left(U_j=1,\, V_j=1\right) + \Pr\left(U_j=0,\, V_j=0\right)\right)^N\\ &= \left(c^2 + (1-c)^2\right)^N = \left(2c^2 - 2c + 1\right)^N\end{aligned}$$

Similarly, the probability that the two connectivity vectors differ from each other by exactly d PNs is:

$$\begin{aligned}\Pr\left(H\left(\vec{U},\vec{V}\right)=d\right) &= \Pr\left(\sum_{j=1}^{N}(U_j - V_j)^2 = d\right) = \left(\Pr\left(U_j=1,\, V_j=1\right) + \ldots\right.\\ &\quad \left. + \Pr\left(U_j=0,\, V_j=0\right)\right)^{N-d} \cdot \left(\Pr\left(U_j=1,\, V_j=0\right) + \ldots\right.\\ &\quad \left. + \Pr\left(U_j=0,\, V_j=1\right)\right)^{d} \cdot \begin{bmatrix} N \\ d \end{bmatrix}\\ &= \left(2c^2 - 2c + 1\right)^{N-d} \cdot \left(2c(1-c)\right)^{d} \cdot \frac{N!}{d!(N-d)!}\end{aligned}$$

This yields a theoretical probability-density function (PDF) for the Hamming distance between connectivity vectors (**Figures 2B–D**). Note that for all values of c the PDFs are always rather narrow (**Figure 2B**), with most of their mass concentrated close to their mean value. This is a key property of binomial distributions with large values of N, and implies that most pairs of connectivity vectors in a system obeying our basic assumptions will differ by similar values, well predicted by their mean Hamming distance. Note also, that the PDF centered on the highest value is for c = ½, the connectivity value measured between PNs and KCs in the locust. **Figures 2C,D** provide a closer look at this particular case (see next section).

Connectivity ½ thus maximizes differences between PN–KC connectivity vectors. I demonstrate this graphically in **Figure 3** using elementary Venn diagrams. Two different KCs, each of which samples PNs randomly and independently with probability c, thus define two sets of PNs (I call these sets u and v). Each large (open) circle in **Figure 3A** represents the entire PN set (with area N), the two smaller circles within it mark the PN subsets u and v sampled by our two KCs (with average area $N \cdot c$ each; the value of c is indicated above each diagram).

The set of PNs sampled by both KCs (the *overlap* of the two PN sets) is the intersection of u and v, the number of PNs it includes on average is

$$\langle \|u \cap v\| \rangle_{u,v} = \left\langle \sum_{j=1}^{N} U_j V_j \right\rangle_{\vec{U},\vec{V}} = Nc^2$$

as demonstrated by the dark-shaded areas in **Figures 3A,B**. Similarly, the set of PNs sampled by *exactly* one of the two KCs (the *non-overlapping portion* of inputs to the two KCs, or their

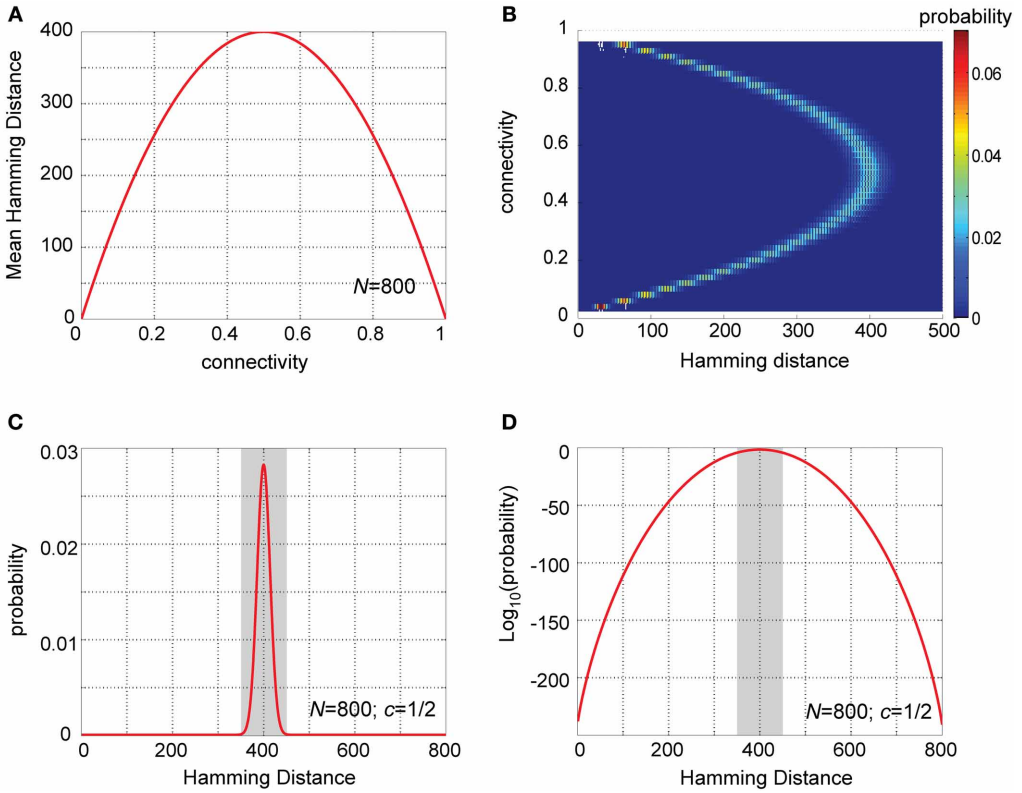

FIGURE 2 | Analysis of Hamming distances between PN–KC connectivity vectors. (A) Analytic solution for mean Hamming distance between PN–KC connectivity vectors as a function of connectivity, c (for $N = 800$ PNs). Hamming distance is a parabola in c with maximum at $c = ½$. **(B)** Probability-density functions (PDFs) of Hamming distances between PN–KC connectivity vectors, $\Pr(H(U, V) = j)$, calculated as a function of connectivity. Each row corresponds to one PDF (with PN–KC connectivity value on ordinate); color represents probability. $N = 800$ PNs assumed in all cases. Note narrow distribution of Hamming distances around their mean for all values of c. Note PDF centered around highest value (400) is for $c = ½$. **(C)** Theoretically calculated PDF of Hamming distances between PN–KC connectivity vectors for parameters measured in the locust olfactory system ($N = 800$, $c = ½$). Note most common Hamming distance value, 400 (as each of two KCs samples on average 200 PNs that the other does not, see text). Shaded area, interval [350,450] in which most probability-density (0.9997) is concentrated. **(D)** Linear-log plot of same PDF as in **(C)**. Shaded area, interval [350,450]. Note miniscule values away from the mean (outside shaded area).

symmetric difference, Δ, in set theory terms) is the union of u and v minus their intersection; the average number of PNs it includes:

$$\langle \|u \Delta v\| \rangle_{u,v} = \langle \|(u \cup v) \backslash (u \cap v)\| \rangle_{u,v}$$
$$= \left\langle \sum_{j=1}^{N} U_j + \sum_{j=1}^{N} V_j - 2 \sum_{j=1}^{N} U_j V_j \right\rangle_{\vec{U},\vec{V}} = 2Nc(1-c)$$

which corresponds to the light-shaded area in **Figures 3A,B**. This tells us how much these two KCs differ on average in PN ensembles they sample (or in their "receptive fields" in terms of input). This area is small when c is very low or very high, and maximal when $c = 0.5$ (as seen in **Figure 3A**, and more clearly in the bar graphs in **Figure 3B**). In fact, this expression is also precisely the result we got for Hamming distance between connectivity vectors (see above and **Figure 2A**). The white areas ("None" in **Figures 3A,B**) correspond to PNs not sampled by either of two KCs. Both the average union and average intersection of the two PN ensembles increase monotonically with connectivity, but the difference between them (the non-overlapping ensemble) peaks at ½ (**Figure 3C**).

Differences between receptive ranges (or "receptive fields") of two target neurons are thus large when they each sample an *intermediate* proportion of the source-population—sampling either a very small or very large proportion yields much smaller non-overlapping ensembles, hence source-populations less different from each other.

PROPERTIES OF THE CONNECTIVITY MATRIX: PLUGGING IN REAL-DATA VALUES

To sense how the above translates into biological reality, let us apply these calculations to the connectivity matrix of the locust olfactory system. For values relevant to the locust ($N = 800$ PNs and $c = ½$) the mean Hamming distance between two PN–KC connectivity vectors is 400; two randomly chosen KCs will thus overlap by 200 connected PNs on average, and each of the KCs will on average sample 200 PNs which the other does not. There will be an additional 200 PNs which are not sampled by either of the two KCs.

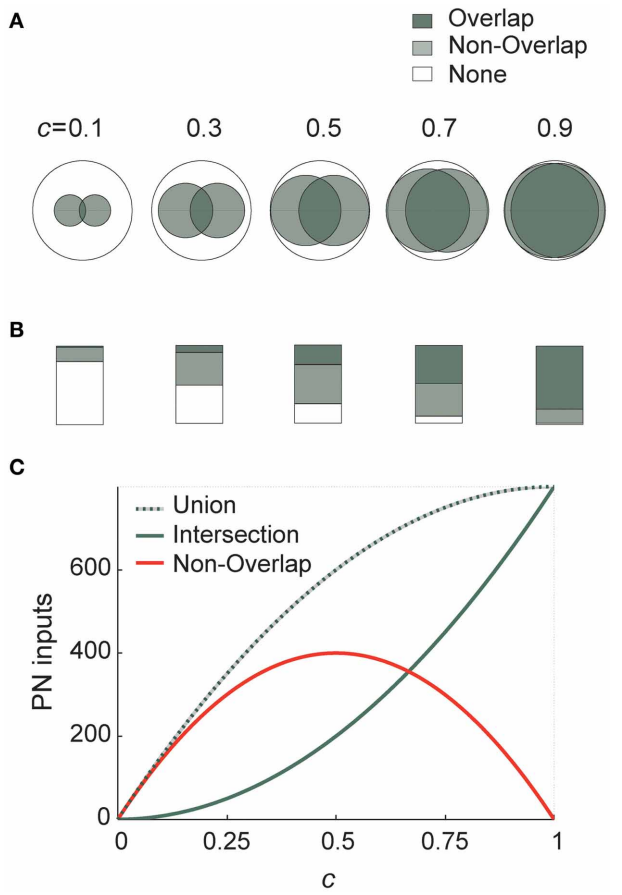

FIGURE 3 | Connection probability ½ maximizes differences between KC input-populations. (A) Schematic representation of two KC inputs using Venn diagrams. Each large (empty) circle represents the entire PN population; the shaded circles within it represent two average KCs receiving connections from subsets of these PNs (with probability indicated above each diagram). Total shaded area (light-shaded + dark-shaded) represents the union of the two KC input-ensembles (or "receptive fields"), while the dark-shaded area alone is their intersection. The light gray area thus corresponds to the non-overlapping portion of the input-ensembles (union minus intersection), or to how different KCs are from each other in terms of input. **(B)** Same as in **(A)** using bar graphs. Each large rectangle represents the entire PN population; shaded areas use same color code and same connection probabilities as in **(A)**. **(C)** Analytically calculated curves of the union (dotted line), intersection (solid gray) and their difference (red) for two KCs in terms of PN input, as a function of PN–KC connection probability c. While the two former are both monotonically increasing, the latter is maximized at $c = ½$. Representations of the outside world are thus spread maximally across the target neuron population for connectivity ½.

Figures 2C,D show the predicted distribution of Hamming distances between PN–KC connectivity vectors in the locust. Note the mean Hamming distance between two KC connectivity vectors (400) is also by far the most common value; it occurs with probability 0.028. The main mass of the distribution is tightly concentrated around the mean value (**Figure 2C**): 0.9997 out of a total mass of 1 of the PDF lies within ±50 PN inputs from the mean (shaded area in **Figure 2C**); randomly chosen pairs of KCs will thus almost always (in 99.97% of cases) have input-ensembles differing by 350–450 PNs. The PDFs take extremely small values further away from the mean, as better seen on a semi-logarithmic scale (**Figure 2D**, shaded area is same interval): note the miniscule probabilities outside the interval 350–450. The probability that two different KCs will sample the exact same PN ensemble is $\sim 10^{-241}$, and the probabilities that their input-ensembles will differ by 1, or 2, or 3 inputs are 10^{-238}, 10^{-235}, and 10^{-233}, respectively—vanishingly small numbers in all these cases.

MODEL RESULTS II: NEURONAL ACTIVITY AND PROPERTIES OF INPUT TO KCs

Up until now, we only considered the properties of the connectivity matrix, $\overleftrightarrow{W}$. To see what happens when neural activity is added in, let us put some flesh on the dry skeleton, and proceed to explore the aggregate input to KCs ($\vec{k}$) during network activity—corresponding to their sub-threshold membrane-potential. The symbol Ψ denotes the mean aggregate input to a KC, averaged over all possible PN-population states and across all KCs. Then

$$\Psi \equiv \langle k_i \rangle_{\vec{S},i} = \left\langle \left\langle \sum_{j=1}^{N} W_{ij} S_j \right\rangle_{\vec{S}} \right\rangle_i = \left\langle \sum_{j=1}^{N} W_{ij} \langle S_j \rangle_{\vec{S}} \right\rangle_i$$
$$= p \cdot \sum_{j=1}^{N} \langle W_{ij} \rangle_i = Npc$$

the mean aggregate input to a KC during our arbitrary time window is thus a simple product of the number of PNs, probability of spiking in a single PN during this snapshot and PN–KC connection probability (**Figure 4A**).

Λ will denote the variance of k_i, averaged across all KCs and over all possible PN-population states (**Figure 4B**) (see Appendix A2 for full derivation):

$$\Lambda \equiv \langle \mathrm{var}(k_i) \rangle_i = \left\langle \left\langle \left(k_i - \langle k_i \rangle_{\vec{S}}\right)^2 \right\rangle_{\vec{S}} \right\rangle_i$$
$$= \left\langle \left\langle \left(\sum_{j=1}^{N} W_{ij} S_j - \Psi \right)^2 \right\rangle_{\vec{S}} \right\rangle_i$$
$$= \sum_{j=1}^{N} \sum_{k=1, j \neq k}^{N} \langle W_{ij} W_{ik} \rangle_i \langle S_j S_k \rangle_{\vec{S}} + \ldots$$
$$+ \sum_{j=1}^{N} \left\langle W_{ij}^2 \right\rangle_i \left\langle S_j^2 \right\rangle_{\vec{S}} - 2\Psi \cdot \sum_{j=1}^{N} \langle W_{ij} \rangle_i \langle S_j \rangle_{\vec{S}} + \Psi^2$$
$$= Npc(1 - pc)$$

So we have explicitly expressed the mean and variance of the aggregate input k_i (**Figures 4A,B**) as a function of basic network parameters. Note that variable k_i is a product of two mutually independent, binomially distributed variables: the momentary vector of spiking in the PN population [a binomial with parameters (N, p)], and the vector of connections between the PN set and the KC [a binomial with parameters (N, c)]. Their dot product, k_i, is also a binomial variable, with parameters N and $p \cdot c$, as indicated by the calculations of Ψ and Λ.

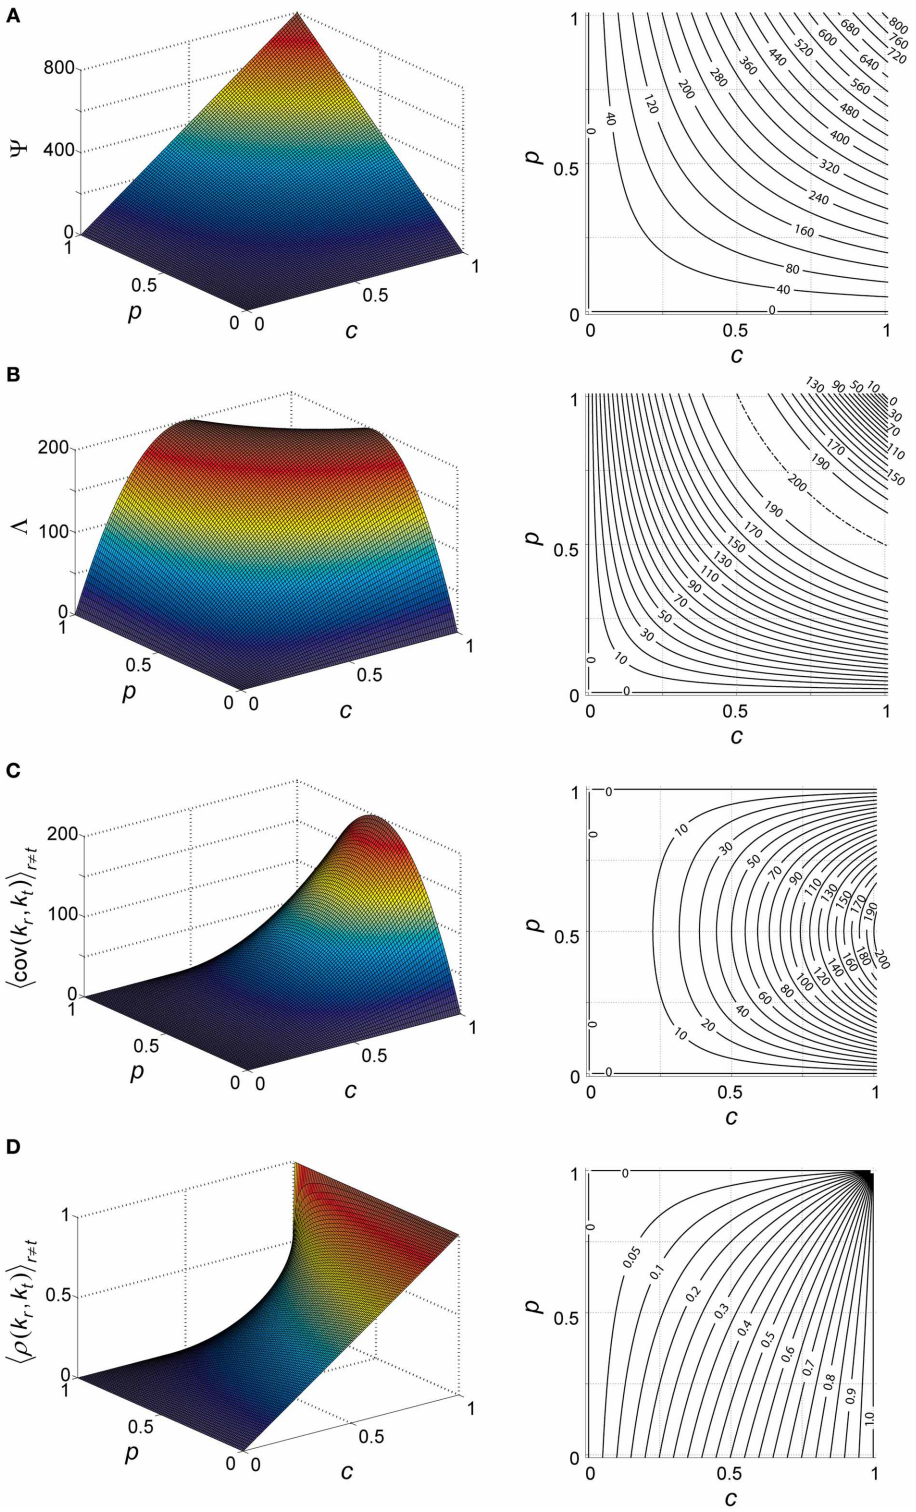

FIGURE 4 | Theoretical properties of network input to KCs. **(A–D)** Analytically calculated properties of the aggregate input to KCs (k) during network activity. Aggregate input is also analogous to KC membrane-potential (see text). Each property plotted as a function of PN-spiking probability p and PN–KC connectivity c, and averaged over all antennal-lobe states. $N = 800$ PNs is assumed in all cases. Left, surface plot; right, contour plot. Contour intervals are identical within each plot. Dash-dot lines indicate ridge contours. For clarity, isoline values are sometimes indicated beside plot (when contour lines are too dense for inline labeling). **(A)** Mean aggregate input per KC, Ψ [units of PNs]; contour interval, 40. **(B)** Variance of aggregate input per KC, Λ [units of PNs]; contour interval, 10. **(C)** Covariance between aggregate inputs to two KC [units of PNs]; contour interval, 10. **(D)** Correlation coefficient between aggregate inputs to two KC [unitless]; contour interval, 0.05.

To what extent are aggregate inputs to two KCs correlated with each other? Calculating their covariance (**Figure 4C**) we get (Appendix A3):

$$\begin{aligned}\langle \mathrm{cov}(k_r, k_t)\rangle_{r\neq t} &= \left\langle\left\langle \left(k_r - \langle k_r\rangle_{\vec{S}}\right)\left(k_t - \langle k_t\rangle_{\vec{S}}\right)\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\ &= \left\langle\left\langle\left(\sum_{i=1}^{N} W_{ri}S_i - \Psi\right)\left(\sum_{j=1}^{N} W_{tj}S_j - \Psi\right)\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\ &= \sum_{i=1}^{N}\sum_{j=1,\, i\neq j}^{N} \langle W_{ri}W_{tj}\rangle_{r\neq t}\langle S_iS_j\rangle_{\vec{S}} + \ldots \\ &\quad + \sum_{i=1}^{N}\langle W_{ri}W_{ti}\rangle_{r\neq t}\langle S_i^2\rangle_{\vec{S}} - \ldots \\ &\quad - \Psi\sum_{i=1}^{N}\langle W_{ri}\rangle_{r\neq t}\langle S_i\rangle_{\vec{S}} - \ldots \\ &\quad - \Psi\sum_{j=1}^{N}\langle W_{tj}\rangle_{r\neq t}\langle S_j\rangle_{\vec{S}} + \Psi^2 = Nc^2p(1-p)\end{aligned}$$

and their correlation coefficient (**Figure 4D**) is:

$$\begin{aligned}\langle \rho(k_r, k_t)\rangle_{r\neq t} &= \left\langle \frac{\mathrm{cov}(k_r, k_t)}{\sqrt{\mathrm{var}(k_r)\cdot\mathrm{var}(k_t)}}\right\rangle_{r\neq t} = \frac{Nc^2p(1-p)}{\sqrt{\left(Npc(1-pc)\right)^2}} \\ &= \frac{c-pc}{1-pc}\end{aligned}$$

Note that both covariance and correlation coefficient have non-negative values in our model (as p and c are probabilities, $1 \geq p$, $c \geq 0$, and N is the number of PNs, $N \geq 0$); this is expected in a network with architecture as described—with all connections feed-forward and excitatory—and with no correlations assumed between external inputs to the system. For $c = 1$, the correlation coefficient is 1 (as all KCs see the exact same input); for $c = ½$ the correlation coefficient is $\frac{1-p}{2-p}$, ranging between 0 and ½.

MODEL RESULTS III: INTER-KC DIFFERENCE IS MAXIMAL FOR CONNECTIVITY ½

We now touch a fundamental question: for given network parameters, how much do target neurons differ from each other in their aggregate inputs? This will tell us how much two KCs differ in sub-threshold membrane-potentials within a given cycle in the active network (earlier we asked how connectivity vectors differ; Section "Model Results I: Exploring Properties of the Connectivity Matrix"). Let us calculate the *difference*, $D(X, Y) \equiv |X - Y|$, between aggregate inputs to two KCs (see Appendix A4 for alternative derivation):

$$\begin{aligned}D(k_r, k_t) &\equiv \left\langle \langle |k_r - k_t| \rangle_{\vec{S}}\right\rangle_{r\neq t} = N\cdot\left\langle\left\langle (W_{ri}S_i - W_{ti}S_i)^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\ &= N\cdot\left\langle\left\langle (W_{ri}S_i)^2 - 2W_{ri}W_{ti}S_i^2 + (W_{ti}S_i)^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\ &= N\cdot\Big(\langle W_{ri}\rangle_{r\neq t}\cdot\langle S_i\rangle_{\vec{S}} - 2\langle W_{ri}W_{ti}\rangle_{r\neq t}\cdot\langle S_i^2\rangle_{\vec{S}} + \ldots \\ &\quad + \langle W_{ti}\rangle_{r\neq t}\cdot\langle S_i\rangle_{\vec{S}}\Big) \\ &= N\cdot\left(cp - 2c^2p + cp\right) = 2pNc\,(1-c)\end{aligned}$$

It is straightforward to see that when taken as a function of PN–KC connection probability, D is maximal for $c = ½$; this holds for any positive p and N (i.e., for all biologically relevant cases, with non-zero PN-firing probability and more than zero PNs in the network). The behavior of D as a function of p and c is shown in **Figure 5A**, and in normalized form in **Figure 5B**.

The above proves that when each target-cell samples half of the source-neurons, the mean difference between inputs to any two targets is maximized. Stated differently, each KC is on average maximally different from all other KCs in the information it carries about the external state.

INTER-KC DIFFERENCE: PLUGGING IN REAL-DATA VALUES

We can now introduce the values measured experimentally in the locust into our model. At baseline, PNs typically fire at ∼2.5–4 Hz (Perez-Orive et al., 2002; Mazor and Laurent, 2005). The relevant integration time window for KCs is the 50 ms odor-induced oscillation cycle (Perez-Orive et al., 2002); even in the lack of oscillations EPSPs in KCs have a time course of several tens of milliseconds (Jortner et al., 2007). This provides a crude estimate of p, the probability of spiking within the relevant time window:

$p \approx 0.125$–0.2 (Perez-Orive et al., 2002; Mazor and Laurent, 2005);
$c \approx 0.5$ (PN–KC connectivity measurements; Jortner et al., 2007);
$N \approx 800$ (axon count in the PN–KC tract; Leitch and Laurent, 1996).

Introducing these numbers into $D = 2pNc\,(1-c)$, the mean difference between two KCs is equivalent to 50–80 PN inputs (**Figure 5C**). If only 100 PNs converged onto each KC ($c = 0.125$), the mean difference would be 22–35 PNs, and with only 10 PNs per KC ($c = 0.0125$, as previously estimated; Perez-Orive et al., 2002), it would be equivalent to only 2.4–4 PNs at baseline (**Figure 5C**)!

During odor presentation, average PN firing-rates do not change significantly over the population (Mazor and Laurent, 2005). However, as PN-spikes are now confined to about half the oscillation cycle (the rising phase; Laurent and Davidowitz, 1994; Laurent et al., 1996; Wehr and Laurent, 1996), p effectively increases by ∼factor 2 (by virtue of the time window "shrinking"). The mean difference between two KCs thus increases to 100–160 PNs during odor; if the fan-in were 100 PNs per KC ($c = 0.125$), or 10 PNs per KC ($c = 0.0125$), the mean difference would become 44–70 PNs, or 5–8 PNs, respectively.

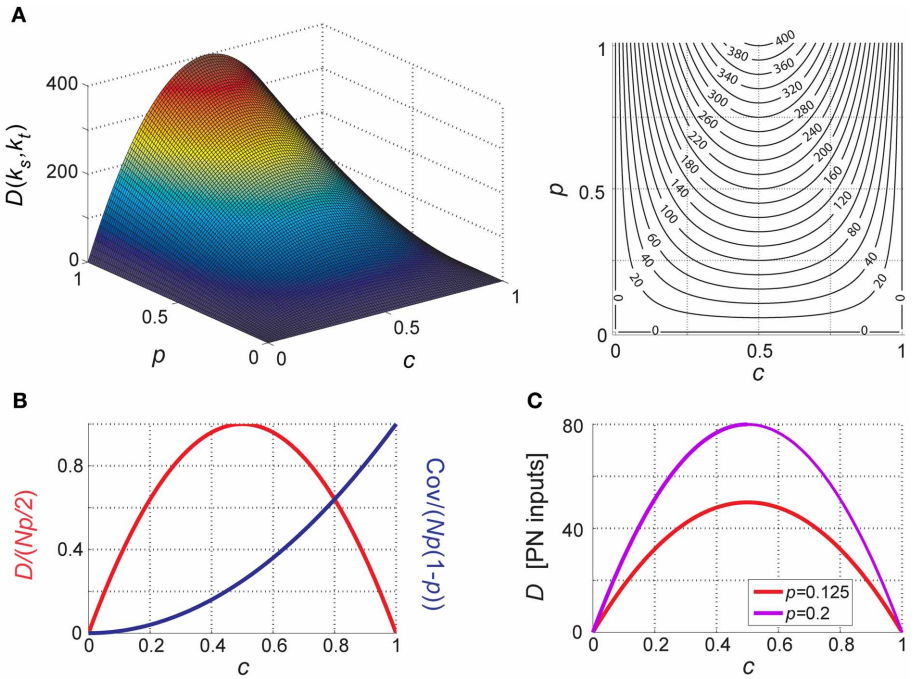

FIGURE 5 | Distance between KC inputs (or membrane-potentials). **(A)** Difference (D) between aggregate inputs to two different KCs (averaged over all pairs of different KCs and over all antennal-lobe states) as a function of PN-firing probability p and PN–KC connection probability c. D is a parabola in c (with maximum at $c = ½$) and increases linearly with p and with N. $N = 800$ PNs is assumed. D is in PN inputs. Left, surface plot; right, contour plot, contour interval is 20. **(B)** Normalized difference and covariance between KCs as a function of connectivity c; normalized D and cov are unitless, independent of N and p and vary between 0 and 1. **(C)** Predicted difference between two KCs for PN-firing parameters measured of in the locust olfactory system (at both extremes of the range): $N = 800$, $p = 0.125$ (2.5 Hz, red), and $p = 0.2$ (4 Hz, magenta).

MODEL RESULTS IV: ESTIMATING FIRING THRESHOLD AND SPARSENESS

The above observations do not yet relate to KC response properties, as we up to now ignored membrane non-linearities and spiking. What happens when we impose a firing threshold, and assume the KC spikes once it is crossed? We now use the assumption of independence across PNs, and the fact that many of them respond to each odor and during each cycle (according to this model $N \cdot p$ per time window, or 100–160 PNs for values $N = 800$, $p = 0.125$–0.2. According to experimental data, 100–150 PNs fire per cycle; Mazor and Laurent, 2005). With these assumptions, we can apply the Central Limit Theorem (CLT) to the summation of inputs onto a KC: we can treat k as a Gaussian random variable, fully defined by its mean (Ψ) and variance (Λ) which we calculated (Section "Model Results II: Neuronal Activity and Properties of Input to KCs"):

$$k_i = \sum_{j=1}^{N} W_{ij} S_j$$

$$k_i \sim \mathrm{Norm}(\Psi, \Lambda) = \mathrm{Norm}(Npc, Npc(1-pc))$$

where Norm(X, Y) stands for a Normal distribution with mean X and variance Y. So for a given threshold f (in units of PN inputs), the probability of the ith KC crossing the threshold (i.e., spiking) is:

$$\Pr(k_i \geq f) = \frac{1}{\sqrt{\Lambda 2\pi}} \int_{f}^{+\infty} e^{\frac{-(x-\Psi)^2}{2\Lambda}} dx$$

$$= 1 - \frac{1}{\sqrt{\Lambda 2\pi}} \int_{-\infty}^{f} e^{\frac{-(x-\Psi)^2}{2\Lambda}} dx = 1 - Q\left(\frac{f-\Psi}{\sqrt{\Lambda}}\right)$$

where $Q(z)$ denotes the Normal cumulative distribution function (CDF) of variable z (**Figure 6A**). If the firing threshold of KCs f is set to:

$$f = \Psi + z \cdot \sqrt{\Lambda}$$

this will result in a known and defined fraction of all KCs—equal to the area of the tail of the Gaussian [given by the CDF, $1 - Q(z)$]—crossing the firing threshold for a given PN-population state. Similarly, a given KC will cross threshold f in response to a known fraction [again, the area of the tail $1 - Q(z)$] of all PN-population states. This function, illustrated in **Figure 6A**, thus links the KC firing threshold with both the sparseness and the selectivity of the mushroom body neural code (as described in the Introduction).

To demonstrate the usage of this function: a hypothetical feed-forward network which satisfies our assumptions (Section "Model Assumptions") has parameters $N = 100$, $p = 0.5$, and $c = 0.2$. If we wish to "design" a population of target neurons with a particular level of sparseness (say we want each responding to 2.3% of external states), we will set their firing thresholds

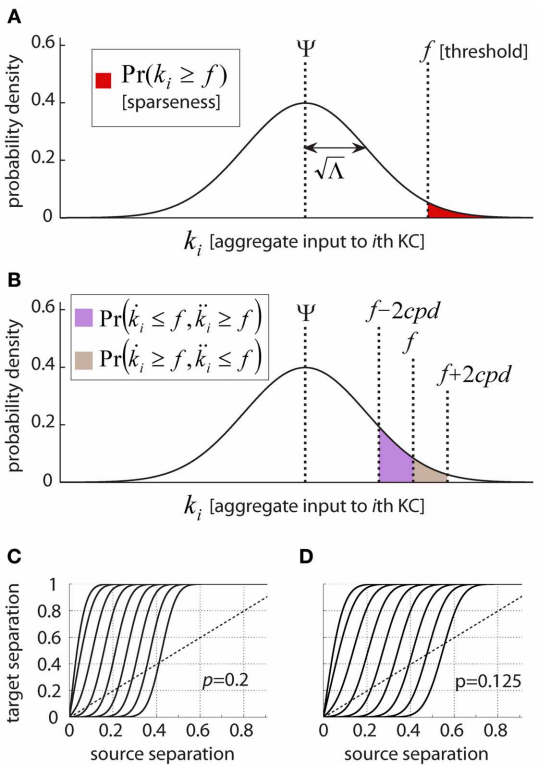

FIGURE 6 | Linking KC firing threshold and response properties. **(A)** Dependence of Kenyon cell response sparseness on firing threshold. Kenyon cell aggregate input, or membrane-potential, k, behaves as a Gaussian random variable with mean Ψ and variance Λ under model assumptions; KC-response probability is the probability of aggregate input crossing firing threshold f (red shaded area). This defines both single-neuron selectivity (i.e., the average fraction of states which evoke firing in a given cell) as well as population sparseness (i.e., the fraction of KCs responding to an average state). **(B)** KC responses to partially overlapping PN activity patterns. Pairs of PN activity vectors differing from each other by exactly d bits present a KC with quantifiably different aggregate inputs. The summed probabilities of the KC responding to the first pattern and not the second (brown) or vice versa (violet) reflect the Hamming distance between KC activity vectors. This function thus expresses the difference between target-states given the difference between source-states. **(C–D)** Difference between Kenyon cell activity vectors as a function of difference between PN activity vectors. Difference is expressed as Hamming distance normalized by the vector-length (or number of neurons in the population). The nine different curves represent different KC firing thresholds, equivalent to (left to right) $[0, 1, 2, \ldots, 8]$ standard deviations above the mean aggregate input. Above the diagonal (stippled line) are regimes where Kenyon cell activity vectors differ from each other more than the corresponding PN activity vectors: i.e., where pattern-separation occurs. Below the diagonal are regimes where different PN activity vectors evoke relatively close KC activity vectors: i.e., generalization/retrieval occurs. Parameters of the locust olfactory system are used, and **(C)** and **(D)** differ by PN-firing rate: **(C)**, $p = 0.2$, corresponding to PNs firing at 4Hz; **(D)**, $p = 0.125$, corresponding to PNs firing at 2.5 Hz; Note how increasing threshold affects curve shape and relative regimes of generalization vs. discrimination.

to be equivalent to 2 standard deviations above their mean aggregate-input (as $1 - Q(2) = 0.023$), or:

$$f = Npc + 2 \cdot \sqrt{Npc(1-pc)} = 10 + 2\sqrt{9} = 16 \quad \text{simultaneous inputs.}$$

A basic, gross prediction which naturally emerges from the threshold-sparseness function is that when the target neurons' threshold is equal to their mean aggregate input (Ψ), each target neuron responds to half of all external states [as $Q(0) = 1 - Q(0) = 0.5$]; this is the most-distributed code possible (also known as a holographic code; Földiák, 2002), and thus sparse coding is conditional on a threshold significantly higher than that, requiring $f >> \Psi$.

THRESHOLD ESTIMATION: PLUGGING IN REAL-DATA VALUES

Let us now test the predictions on firing threshold and sparseness using experimental parameters from the locust. Section "Inter-KC Difference: Plugging in Real-Data Values" shows the network parameters for PN-firing rates ($p \approx 0.125$–0.2), PN–KC connectivity ($c \approx 0.5$) and PN number ($N \approx 800$). Given these, the aggregate input a KC gets is on average

$$\Psi = Npc \approx 50\text{–}80 \text{ PN inputs per cycle}$$

and its standard deviation is:

$$\sqrt{\Lambda} = \sqrt{Npc(1-pc)} \approx 7\text{–}8 \text{ PN inputs per cycle}$$

So for KCs to each respond to ∼1% of all PN-population states, their threshold has to be ∼2.5 SDs above the mean aggregate input: $f \approx 98$–101 PN-inputs for firing rate of 0.2 PN-spikes/cycle (4 Hz), or $f \approx 67$–69 PN inputs for 0.125 PN-spikes per cycle (2.5 Hz).

This result is in good agreement with experimental measurements of KC-response sparseness (∼1–2% using intracellular recordings; Jortner, 2009) and their firing threshold ($f \approx 100$ assuming linear summation and full PN-synchrony; Jortner et al., 2007). The estimate's deviation toward the higher end of the predicted range may be due to supra-linear summation in KCs at depolarized membrane-potentials (Laurent and Naraghi, 1994; Perez-Orive et al., 2002, 2004).

MODEL RESULTS V: NOISE TOLERANCE AND GENERALIZATION

Olfactory stimuli are by nature noisy and variable. PNs show significant trial-to-trial variability when presented with the same odor repeatedly, yet in KCs noise is considerably reduced. Another issue is that some stimuli are similar to each other (either because of chemically related odorant molecules, or because they are mixtures with overlapping components) and others are different. Both points—the way the system tolerates noise and the way it encodes similar or different inputs—are closely related, in that both require us to examine how two overlapping PN-firing patterns are transformed into KC firing patterns.

Recall $\mathbb{S}$, the set of all possible activity vectors of the source-population. Let us define $\{\vec{\dot{S}}, \vec{\ddot{S}}\}$ as the subset of vector-pairs in $\mathbb{S}$ which differ from each other by exactly d-bits. Formally then:

$$\{\vec{\dot{S}}, \vec{\ddot{S}} \in \mathbb{S} | H(\vec{\dot{S}}, \vec{\ddot{S}}) = d\}$$

The aggregate-input vectors to the target-neuron population which are evoked by $\vec{\dot{S}}, \vec{\ddot{S}}$ will be $\vec{\dot{k}}, \vec{\ddot{k}}$, respectively; vectors $\vec{\dot{K}}$ and $\vec{\ddot{K}}$ will be the resulting activity vectors of the

target-population. Aggregate inputs which a single KC gets in response to $\vec{\dot{S}}, \vec{\ddot{S}}$ will differ by:

$$\begin{aligned}\left\langle D(\dot{k}_r, \ddot{k}_r)\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} &= \left\langle \left|\dot{k}_r - \ddot{k}_r\right|\right\rangle_{r,\{\vec{\dot{S}}, \vec{\ddot{S}}\}} \\ &= N \cdot \left\langle \left(W_{rj}\dot{S}_j - W_{rj}\ddot{S}_j\right)^2\right\rangle_{r,\{\vec{\dot{S}}, \vec{\ddot{S}}\}} \\ &= N \cdot \left(\left\langle W_{rj}\right\rangle_r \cdot \left\langle \dot{S}_j\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} - 2\left\langle W_{rj}\right\rangle_r \cdot \left\langle \dot{S}_j\ddot{S}_j\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} + \ldots \right. \\ &\quad \left. + \left\langle W_{rj}\right\rangle_r \cdot \left\langle S_j\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}}\right) \\ &= Nc \cdot \left(\left\langle \dot{S}_j\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} - 2\left\langle \dot{S}_j\ddot{S}_j\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} + \left\langle S_j\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}}\right) \\ &= Nc\left(p - 2p\left(1 - \tfrac{d}{N}\right) + p\right) = 2cpd\end{aligned}$$

and so

$$\ddot{k}_r = \dot{k}_r \mp 2cpd$$

Earlier we linked KC aggregate input and firing threshold to their firing probability (Section "Model Results IV: Estimating Firing Threshold and Sparseness"; **Figure 6A**); let us use the same formalism now. The probability that a single KC responds differently to two PN-states is simply the probability one of these states drives it across the threshold and the other does not. This is demonstrated graphically in **Figure 6B** and is exactly the rationale behind the following calculation:

$$\begin{aligned}\left\langle D(\dot{K}_i, \ddot{K}_i)\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} &= \left\langle \left|\dot{K}_i - \ddot{K}_i\right|\right\rangle_{\{\vec{\dot{S}}, \vec{\ddot{S}}\}} \\ &= \Pr\left(\dot{K}_i = 1, \ddot{K}_i = 0\right) + \Pr\left(\dot{K}_i = 0, \ddot{K}_i = 1\right) \\ &= \Pr\left(\dot{k}_i \geq f, \ddot{k}_i < f\right) + \Pr\left(\dot{k}_i < f, \ddot{k}_i \geq f\right) \\ &= \Pr\left(\dot{k}_i \geq f, \dot{k}_i < f + 2cpd\right) + \ldots \\ &\quad + \Pr\left(\dot{k}_i < f, \dot{k}_i \geq f - 2cpd\right) \\ &= \tfrac{1}{\sqrt{\Lambda 2\pi}} \int_f^{f+2cpd} e^{\frac{-(x-\Psi)^2}{2\Lambda}} dx + \ldots \\ &\quad + \tfrac{1}{\sqrt{\Lambda 2\pi}} \int_{f-2cpd}^{f} e^{\frac{-(x-\Psi)^2}{2\Lambda}} dx \\ &= Q\left(\tfrac{f + 2cpd - \Psi}{\sqrt{\Lambda}}\right) - Q\left(\tfrac{f - 2cpd - \Psi}{\sqrt{\Lambda}}\right)\end{aligned}$$

What about the activity vectors for the KC population, given similar PN input? The mean Hamming distance between two KC activity-patterns given Hamming distance d between the PN activity patterns will simply be the above expression multiplied by the number of KCs, M. We can thus write:

$$\begin{aligned}&\left\langle H(\vec{\dot{K}}, \vec{\ddot{K}}) \,\middle|\, H(\vec{\dot{S}}, \vec{\ddot{S}}) = d\right\rangle \\ &\quad = M\left(Q\left(\tfrac{f + 2cpd - \Psi}{\sqrt{\Lambda}}\right) - Q\left(\tfrac{f - 2cpd - \Psi}{\sqrt{\Lambda}}\right)\right)\end{aligned}$$

NOISE TOLERANCE AND GENERALIZATION: PLUGGING IN REAL-DATA VALUES

So how well does the locust olfactory system tolerate noise? The results are shown in **Figures 6C,D**, where I feed into the relation derived in the previous section the parameters from the locust circuitry. Two PN activity-patterns, differing by 0–800 bits (x axis, normalized to 0–1) evoke KC activity-patterns differing by 0–50,000 bits (y axis, normalized to 0–1). The diagonal (stippled lines) in both figures shows where the hypothetical curve would pass if normalized distance between representations would not change in transition from PNs to KCs. In fact, the relation has a sigmoid shape, meaning PN patterns close to each other become even closer in the KC population; whereas PN patterns which are different become more different (note that distances are normalized to the population size). The nine different sigmoid curves show the relation between input- and output-overlap when the firing threshold is varied (left to right: 0–8 standard deviations above the mean aggregate input). The setting of the firing threshold clearly controls the boundary between generalization and discrimination; a boundary which is surprisingly sharp.

In the locust olfactory system, the KC threshold is located ∼2.5 SDs above the mean aggregate input (commensurate with a sparseness of ∼1% as observed; see Section "Threshold Estimation: Plugging in Real-Data Values"). As seen in **Figures 6C,D**, for this value the Kenyon cell population generalizes (or, tolerates noise) for PN patterns which are within up to ∼50–100 bits away from the PN–KC connectivity vectors, and discriminates for ones which are farther. This means, that with parameters from the locust, the boundary between discrimination and generalization lies in a biologically realistic regime for highly sparse coding (recall that for binary 800–dimensional vectors, over 99.9% of space is removed 350–400 bits from any given vector; **Figure 2**).

SUMMARY OF MODEL RESULTS AND PREDICTIONS

This analytic model produces several insights and predictions, applicable to both the locust olfactory circuitry as well as to feed-forward systems in general. As the model was designed with generality in mind, its results depend only minimally on particular parameter values. Here is a brief summary:

(1) In a feed-forward system with random connectivity, pairs of connectivity vectors from source- to target-population have a maximal Hamming distance for connection probability ½.
(2) Hamming distances between connectivity vectors are mostly very similar to each other and to their mean value; connectivity vectors significantly more similar to each other (or, more different from each other) will be extremely rare (negligible).
(3) Differences in aggregate input (or sub-threshold membrane-potential) between target neurons are maximal for $c = ½$. In the locust antennal-lobe–mushroom-body circuit, where such connectivity is realized, pairs of KCs thus differ from each other by an equivalent of 50–80 PN-inputs during baseline, and of 100–160 PN-inputs when odor is presented. These differences decrease significantly when connectivity shifts away from ½ (in either direction): with connectivity of 10 PNs per KC (as previously estimated;

Perez-Orive et al., 2002) differences between KCs would be equivalent to only 2.4–4 PN inputs (∼factor 20 lower than for $c = ½$).

(4) The standard deviation of sub-threshold membrane-potential in target neurons is maximal when the product of spiking probability in the source-neurons (p) and connectivity between the source- and target-populations (c) is ½. In locust KCs, the standard deviation of membrane-potential is predicted to be equivalent to the sum of 7–8 PN-inputs, or ∼0.6–0.7 mV, in good agreement with experimental measurements.

(5) The covariance of aggregate inputs to two different target neurons will be maximal for $p = ½$, and will increase as $\sim c^2$.

(6) Both the covariance and correlation coefficient between target neurons are predicted to be always positive under the assumptions taken. This is intuitive, given that no correlations were assumed in the external input driving the system, and only feed-forward excitatory connections exist.

(7) The correlation coefficient between target neuron membrane-potentials is expected to range within 0–0.5 for $c = ½$. Particularly, in the locust, where $c = ½$ and PN-spiking probability is 0.125–0.2 per cycle (2.5–4 Hz), correlation coefficients between KCs are predicted to be 0.4–0.5. This remains to be tested experimentally with dual-intracellular KC recordings. A related test—namely measurements of correlations between single-KC membrane-potentials and local-field-potentials—yielded correlation coefficients around 0.3 (Jortner et al., 2007).

(8) The response probability (and sparseness) of target neurons in a feed-forward system with parameters N, p, c is determined by their firing threshold, and is well approximated by the area of a Gaussian tail. The threshold-sparseness function predicts the fraction of states a target cell responds to, and the fraction of target cells responding to any given state. It generates the basic prediction that for a threshold equivalent to the target neurons' mean aggregate input, Ψ (a product of source-neuron firing rate, source-neuron number and connectivity), target neurons will respond to ½ of all source-population states; so to produce sparse coding the threshold must exceed that value: $f >> \Psi$.

(9) Applying the threshold-sparseness function to the locust olfactory system, the firing threshold measured (∼100 inputs, assuming perfect synchrony and linear summation) well predicts the measured KC sparseness level (∼1–2%) and vice-versa (1% sparseness predicts a threshold of ∼70–100 inputs, depending on PN-firing rate).

(10) Given a network with parameters N, p, and $c = ½$, if target neurons have a firing threshold of $f(z) = \frac{Np + z \cdot \sqrt{Np(2-p)}}{2}$ (see Appendix A5), then each target neuron will respond to $1 - Q(z)$ of source-population states, and different target neurons will respond to maximally different states. The difference between the target neurons will on average be $\frac{N \cdot p}{2}$. Combined with adaptive gain control to ensure that f is changed appropriately when p changes (Papadopoulou et al., 2011), this yields a simple way to design a network with an arbitrarily sparse level of activity, and with specific and reliable neural responses to external states.

DISCUSSION: LINKING NETWORK ARCHITECTURE AND NEURAL CODING IN THE ANTENNAL-LOBE–MUSHROOM-BODY CIRCUIT

Integrating theory and experiment, I here discuss how the architecture of the locust olfactory system gives rise to Kenyon cell coding properties: specificity, reliability, low firing rates, correlations, and sparseness, and how these can be utilized to build higher-level representations of the animal's world. Several predictions with potentially broader implications will follow.

CONNECTIVITY ½ MAXIMIZES DIFFERENCES BETWEEN TARGET NEURONS

The key experimental finding motivating this study was that each KC in the mushroom body receives synaptic connections from antennal-lobe PNs with probability ½, each thus sampling 400 of the 800 PNs (Jortner et al., 2007). At first, this result may seem very surprising—because it seems counterintuitive that KC specificity could arise from such broad PN inputs. It makes sense, however, when viewed from a combinatorics perspective: the number of ways to pick n elements out of N is given by the binomial coefficient:

$$\begin{bmatrix} N \\ n \end{bmatrix} = \frac{N!}{n!(N-n)!}$$

This expression is maximal for $n = N/2$, decreasing sharply and symmetrically around it. The fundamental realization that choosing half the elements maximizes the number of possible combinations has dawned independently on several thinkers throughout history—from Pingala (India, 2nd–5th century BC, commentated by Halayudha, 10th century AD), through Al Karaji (Persia, 953–1029), Omar Khayyam (Persia, 1048–1131), Yang Hui (China, 1238–1298), Niccolo Tartaglia (Italy, 1500–1557) to Blaise Pascal (France, 1655).

How is this relevant to the olfactory system? Think of each KC as if picking the PNs it will listen to. If each KC sampled only $n = 1$ of $N = 800$ PNs, there would be exactly 800 ways to pick which PN to sample (similarly, if each KC sampled 799 out of the 800 PNs; where there would be exactly 800 ways to pick which PN *not* to sample). However, when sampling half the PNs, $n = 400$, the number of ways to do so is maximal, and equals $800!/400!400! \approx 10^{240}$. This is an immense number—beyond astronomical—and way too large for any example from nature to demonstrate it. It is roughly equivalent to the number of atoms in the known universe (estimated at $\sim 10^{79}$) taken to the power of three. . .

But as there are only 50,000 KCs in the locust mushroom body, only $5 \cdot 10^4$ combinations are realized out of this vast pool of possibilities. What is the probability that two randomly chosen KCs sample the exact same PN-ensemble? The answer is $\approx 10^{-240}$, which is for all practical purposes zero. And what is the probability that two KCs sample very similar PN-ensembles—that is, ensembles differing from each other by just one, or 2 or 3 inputs? The answers are 10^{-238}, 10^{-235}, and 10^{-233}, respectively—all vanishingly small. In fact, the average pair of KCs will differ by ∼400 PN inputs (**Figure 2C**), which also constitutes the most common case (occurring with probability 0.028), and 99.97% of KC pairs will deviate from it by less than 50 inputs

(**Figures 2C,D**). This stems from a key property of binomial distributions with large N: most of their mass occupies a very narrow band around their mean.

By this reasoning (proven for generalized cases in Sections "Model Results I: Exploring Properties of the Connectivity Matrix" and "Model Results III: Inter-KC Difference is Maximal for Connectivity ½" for connection vectors and membrane-potentials, respectively) each target neuron receives a unique set of source-neuron inputs, very different from that of all other target neurons. KCs are *maximally* different from each other in what they tell us about the world of inputs, because their connectivity vectors are drawn from a pool which is maximal. This feature of the KC population results directly from combinatorics, and from the probability ½ of receiving connections from their source-neurons (**Figure 3**).

A critical comment raised by several colleagues against the above argument is that while this architecture indeed maximizes input separation, this optimum cannot reflect on biological reality. The brain, they argue, could not come so close to it, because the numbers in question are too large to be distinguished from each other by a biological system. In other words, realizing 50,000 combinations of "only" 10^{22} (the number of ways to pick 10 PNs from 800, corresponding to connection probability $c = 0.0125$) would be already immensely sparse; and for all practical purposes 10^{240} (the number of ways to pick 400 from 800) is not "sparser." Moreover, since the mathematical optimum is not necessary, evolution of such connectivity couldn't have possibly been guided by biological selection pressures.

The results presented here (Section "Inter-KC Difference: Plugging in Real-Data Values"; **Figure 5**) refute this criticism. While indeed the binomial coefficient $\begin{bmatrix} 800 \\ m \end{bmatrix}$ rises very steeply with m and soon produces vast numbers, these numbers directly translate into state-dependent differences in aggregate input, or membrane-potential, produced across KCs (**Figure 5B**, normalized difference; **Figure 5C** difference in inputs). If 80 PNs were to connect to each KC (corresponding to $c = 0.1$), the amount by which aggregate inputs to different KCs would differ—and the system's ability to discriminate between external states—would drop approximately to a third of the optimum, and for 10 PNs per KC ($c = 0.0125$) it would drop by a factor of 20. When translated to membrane-potential differences between KCs, this may be critical for readout, especially in the presence of noise. This maximum is thus likely to be meaningful after all, and may account for the exquisitely clean performance of sparse neural systems feeding on noisy input.

WHAT DETERMINES HOW SPARSE THE CODE WILL BE?

KC aggregate inputs differ maximally as a result of the PN–KC connectivity; yet while this property paves the road toward sparse coding, it does not in itself suffice to explain the KCs' rare firing: it is eventually their firing threshold which determines firing probability and response sparseness (Section "Model Results IV: Estimating Firing Threshold and Sparseness"). With such high convergence ratio (400:1), target cells can afford to have a very high firing threshold, which can account for KC specificity, reliability, and low firing rates. Experimental measurements show that KC firing threshold is equivalent to simultaneous activation of ∼100 PN inputs assuming linear summation (Jortner et al., 2007). An estimate based on intracellular recordings (and thus less biased than extracellular studies, as it also captures cells firing rarely or not at all) suggests KCs respond to 1–2% of odors tested (Jortner, 2009). Here, a theoretical function was derived which links firing threshold and response probability (Section "Model Results IV: Estimating Firing Threshold and Sparseness"): it closely predicts the experimental results, estimating the firing threshold necessary to achieve ∼1% KC sparseness at ∼67–101 inputs, depending on PN-firing rates (Section "Threshold Estimation: Plugging in Real-data Values").

The threshold-sparseness function is quite sensitive to activity levels of the input network (Huerta et al., 2004; Jortner et al., 2007; Nowotny, 2009). Since PN population activity produces a range (100–150) of spikes per cycle (Mazor and Laurent, 2005), this can result in instability of the code—causing some external states to activate a large number of KCs and others to activate none at all (Papadopoulou et al., 2011). This requires adaptive gain control of the KC firing threshold to fit the actual activity level of the input; one mechanism shown to maintain output sparseness over a wide range of input conditions in the locust takes place via a large, non-spiking GABAergic interneuron with extensive connectivity and graded release properties; it forms a negative-feedback loop onto KCs and adaptively regulates their population output on a cycle-to-cycle basis (Leitch and Laurent, 1996; Papadopoulou et al., 2011).

In a theoretical exploration of the hippocampus, O'Reilly and McClelland (1994) also find that a "floating threshold" (as they phrase it) is highly useful for determining response sparseness under different input conditions and postulate that adjustment of the threshold can be useful for shifting between pattern-separation (or discrimination, or new-category formation) and pattern-completion (or generalization, or recall).

It should be noted that the threshold-sparseness function derived here is independent from the results on connectivity, and it can be applied to systems with any parameters.

EFFECTS OF PN–KC CONVERGENCE: RELIABILITY AND CORRELATIONS

While overall PN–KC circuit-architecture is highly divergent due to the increase in dimensionality, the connectivity scheme makes single KCs receive massively convergent input (from 400 PNs). Together with the high and adaptive KC-threshold, this convergence sub-serves the KCs' reliable, low-noise performance: summing many PN-inputs prior to KC threshold (f) crossing is equivalent to massive averaging of PN activity. This reduces the significant variability (i.e., noise) present in individual, cycle-wise PN responses by a factor of $1/\sqrt{f}$; in locust KCs, where $f \approx 100$, noise is thus reduced ∼10-fold.

Another interesting effect of this convergence is the coexistence of correlations and differences in the mushroom body. While membrane potential-differences between different KCs are maximized, they are still predicted by the model to co-vary significantly (**Figure 5B**), with correlation coefficients of 0.4–0.5 [see Section "Summary of Model Results and Predictions"; Prediction (7)]. Indeed, while the mushroom body code is highly sparse and specific (Perez-Orive et al., 2002; Jortner, 2009), a

salient property of KC intracellular membrane potentials is their strong correlations with the mushroom body local field potential (Laurent and Naraghi, 1994; Jortner et al., 2007), implying that they are also highly correlated with each other. How can strong correlations between cells, which we naturally tend to associate with similarity, exist side by side with maximal difference between them?

To answer this apparent paradox we examine the inputs KCs receive vs. the outputs they produce. Correlations between membrane-potentials of two KCs result from massive overlap in their aggregate input: they share on average ∼200 incoming PN-synapses, and 25–40 active PN-inputs per oscillation cycle. The relevant feature for the system is, however, the number of inputs by which they do *not* overlap (**Figure 3**): each also receives on average ∼200 PN synapses (25–40 active ones per cycle) which the other does not; so they differ by 400 synapses (50–80 active inputs per cycle). The non-linearity imposed by the KC-threshold makes the two properties—correlations and difference—strongly diverge at this point: two KCs can get highly correlated inputs, yet may easily sit across different sides of the threshold, which in turn determines who will fire and who won't; both correlations and differences can thus coexist between them.

The general message is that while sub-threshold correlations naturally arise from input overlap in highly interconnected systems, they do not necessarily imply similarity in function (or output) between neurons; depending on network design and on the parameter taken as readout, the non-overlapping input may outweigh the overlap (as shown for KCs). Eventually, the non-linearity of thresholding enables brains to parse the world into percepts and build representations from them. Membrane-potential correlations between KCs may in this case be side effects of the interconnected architecture, rather than a computational feature of the code.

NEURAL DESIGN-PRINCIPLES FOR GENERATING A SPARSE CODE

As pointed out in the Introduction, a prerequisite for understanding a neural system is characterizing its basic features—individual components, connectivity and external input. Formulating higher properties in terms of these features bridges levels of description and thus constitutes deeper understanding. This approach was used here link network design and sparse coding in the antennal-lobe–mushroom-body circuit. The experimentally measured parameters *f, c, p*—corresponding exactly to the individual unit input–output function, connectivity and input—were used to express distance-measures between connectivity vectors and between target neurons, sub-threshold behavior and coding sparseness.

Three main principles govern the design of the antennal-lobe–mushroom-body circuit: First, there are many more target cells than source cells (∼50,000 vs. ∼800); a factor $\sim 10^2$ increase in dimensionality between the two odor-representations. Second, the probability of connection between the principal neurons of both relays is ∼½; each target thus samples ∼400 of 800 sources. Third, target-cell firing threshold is high, equivalent to simultaneously activating ∼100 of their inputs, and can be fine-tuned to fit different activity-levels of the source network.

Due to the high threshold, each external state (or here, PN activity pattern) activates only a small subset of KCs. However, due to the connectivity scheme, different external states activate different KC-subsets. The activation of very small, very different subsets of cells in response to different external states suffices to produce a sparse and selective neural code as we defined it (see "Introduction", and Jortner et al., 2007). At the same time, the high PN convergence onto individual KCs explains why KCs are so reliable on the one hand, and on the other hand why their membrane-potentials are noticeably correlated with the local-field potential (Laurent and Naraghi, 1994; Perez-Orive et al., 2002; Jortner et al., 2007), an observation that initially seemed paradoxical (Jortner et al., 2007). Finally, with thresholds so high, it is not surprising that the chances of "accidental" spiking are very small, and that KC spontaneous-firing rates are extremely low.

The design principles described here thus lead to reliable, specific—sparse as well as selective—representations of random olfactory-percepts in the mushroom body, and form a simple way to make a sparse code spontaneously emerge, with no need for a "guiding hand" such as learning or predetermined connections.

The total number of KCs, their fraction activated per external state, the levels of noise, and the cost function of classification errors will together determine how state-space is tiled—or how many odors can be reliably encoded by the mushroom bodies (and also, how many KCs are needed to encode a certain number of odors). A meaningful estimate is beyond the scope of this work, as a critical parameter—how distant KC representations must be from each other within noise constraints—is unknown.

Directly from the above principles emerges a simple recipe for designing networks with optimal separation of representations and arbitrarily specific responses. If two neuronal populations have feed-forward connectivity with probability ½, N source-neurons firing p spikes per characteristic time, and target neuron firing threshold is equivalent to $f(z) = \frac{Np + z \cdot \sqrt{Np(2-p)}}{2}$ (Appendix A5), then each target neuron will respond to a known proportion of source-population states (the area under Gaussian tail $1 - Q(z)$), and different neurons will respond to maximally different states. Adaptive gain control should be implemented to ensure f changes appropriately when p changes (Papadopoulou et al., 2011). At any given time, target neurons' aggregate inputs (momentary membrane-potentials) will on average differ by $\frac{N \cdot p}{2}$.

GENERALIZATION vs. DISCRIMINATION

An inherent dilemma when parsing sensory input is where to draw category-lines. Sometimes a stimulus must be recognized—i.e., grouped into an already-existing category—even if it has never been previously encountered in the exact same form. This allows recognition of sensory stimuli in the presence of noise, as well as grouping things together into meaningful categories (i.e., generalization), both of which are essential requirements for the brain to perform its tasks.

In other cases, stimuli which may be very close to each other need to be told apart. Discrimination is critical when selecting food, for example. An extreme case is when the system needs to decide that something is completely novel and merits a new category of its own.

It is important to recognize that these tasks—discrimination and generalization—contradict each other to some extent, yet

sensory systems need to be able to do both, and sometimes on the very same stimulus: something smells like a fruit (generalization), but clearly does not smell like an apple, though (discrimination).

The model presented here provides some intuition on how this may happen. As shown in **Figures 6C,D**, the same network can perform both tasks: with the sigmoid-shaped relation between source- and target-separation, stimuli close to each other at the source-layer will be generalized by the population, whereas stimuli farther from each other will be discriminated. The boundary between discrimination and generalization is rather sharp; and its location is determined by the firing threshold, which can be adapted (Papadopoulou et al., 2011).

KENYON CELLS CAN SERVE AS BUILDING BLOCKS FOR MEANINGFUL (AND PLASTIC) REPRESENTATIONS AT THEIR TARGETS

The basic question we began our journey with is how the brain creates specific, high-level, and eventually ecologically meaningful percepts. The antennal-lobe–mushroom-body transformation described above achieves a major step in this direction by separating representations and giving rise to specific and random responses. However, it remains to be discussed how these random response properties lead to ecologically relevant percepts, and how this fits into the mushroom body's widely accepted role in learning [reviewed in Heisenberg (1998)].

The distribution of connection strengths between PNs and KCs is rather narrow (Jortner et al., 2007); in addition PN–KC synapses show no short-term plasticity, such as homo- or heterosynaptic facilitation or depression (Jortner et al., 2007). While these observations do not rule plasticity out altogether, they definitely do not support plasticity playing a key role at PN–KC synapses.

What happens at the transformation to the next relay? Dendritic trees of β-lobe neurons (one of the main classes of mushroom body outputs) are planar and oriented perpendicular to the KC-axon tract; this structure suggests that β-lobe neurons can integrate precisely timed neural activity over a potentially wide subpopulation of KCs (Li and Strausfeld, 1997; MacLeod et al., 1998). Cassenaer and Laurent (2007) showed that connectivity from KCs to β-lobe neurons is low (∼2%), individual active synapses are relatively strong (1.58 ± 1.11 mV) and exhibit salient spike-timing dependent plasticity, which is sensitive even to single action-potentials.

It is thus attractive to imagine the transformation of information from the antennal-lobe to the mushroom body as happening via widespread, random (or partially random) and largely fixed connections—designed to spread neuronal information optimally and create discrete, specific and reliable representations of random features. This would prepare it for further computation in downstream areas, such as the β-lobe—where more complex ideas can then be constructed from these elementary building blocks, much like words and phrases are constructed from an alphabet (Barlow, 1972; Stryker, 1992). Hence as different KCs respond specifically to various and different chemicals (or classes of chemicals), proper wiring of connections and selection/tuning of their strengths can generate high-level, invariant and "meaningful" representations. For example, a hypothetical downstream neuron responding only to odors associated with locust foods could easily be constructed by connecting onto it only KCs firing in response to various 5- and 6-carbon chained alcohols, aldehydes, and esters which are common odorants in grassy plants (cheerfully nicknamed "green odors"; Hopkins and Young, 1990; Bernays and Chapman, 1994). Similarly, some downstream neurons can respond to odors indicating plant toxicity (for examples of such chemical cues see Cottee et al., 1988).

At this downstream stage, learning (i.e., tweaking of incoming synapses from KCs) can shape and tune these representations, molding them to the animal's specific surroundings. Locusts, as many other generalist animals, readily adapt their food-preferences to seasonal- and regional-variation of plants, their nutritional value and the animal's needs (see Cooper-Driver et al., 1977; Bernays et al., 1992; Bernays and Chapman, 1994), and learning plays an important role in this (Dukas and Bernays, 2000; Behmer et al., 2005). It is likely, that learning a different food-preference is accomplished by changes at KC–β-lobe synapses, based on positive- and/or negative-reinforcement signals—originating, for example, from the digestive system (Behmer et al., 1999) and relayed via neuromodulatory reward/punishment signals, as shown in a variety of insect species (Hammer and Menzel, 1995, 1998; Schwaerzel et al., 2003; Unoki et al., 2005).

Learning can thus sculpt and tune higher neuronal representations, bringing neurons downstream of the mushroom body to respond to "meaningful" stimuli; i.e., stimuli with ecological importance for the animal (Barlow, 1972, 1985). Wiring each of these β-lobe neurons further, to directly trigger a relevant motor-program (e.g., for eating, avoidance, escape, etc.) would close the loop from perception to action. This would result in a simple neural system which receives complex, high-dimensional and noisy input and produces reliable animal behavior in response to it—in other words, a simple brain that works—and where we are approaching a deeper mechanistic understanding of the process.

SWITCHING BETWEEN CODING SCHEMES

A large body of work is focused on sparse codes, pointing out their many benefits (e.g., Barlow, 1972; Palm, 1980; Baum et al., 1988; Kanerva, 1988, 1993; Földiák, 1990, 2002). Sparse codes are attractively easy to read out, as few spikes from few neurons translate to meaning, eliminating the need to integrate over an entire population or over a long time (Földiák, 1990, 2002). Forming associations is easy, as learning has to act on few nodes; in the theoretical limit-case (one-neuron-per-percept) tuning a single synapse suffices to (asymmetrically) link two percepts (Palm, 1980; Baum et al., 1988; Kanerva, 1993). Complex, meaningful ideas can be constructed by wiring-together basic random percepts (see Section "Kenyon Cells Can Serve as Building Blocks for Meaningful (and Plastic) Representations at their Targets"). Finally, sparse codes are metabolically economical, although an energetic trade-off exists between firing few spikes and maintaining many cells (Levy and Baxter, 1996; Attwell and Laughlin, 2001). Sparse codes are thus attractive and economical substrates for computation.

On the other hand, sparse coding has several serious drawbacks. It is wasteful in hardware, as each neuron participates only in a small fraction of representations (each percept requires

devoted neurons and representations rarely share the same cells). It is also extremely sensitive to neuronal damage, as losing neurons results in loss of precepts or memories (Földiák, 2002).

Sparse codes thus seem unlikely candidates for applications such as long-term memory storage, but they are very well suited for applications such as short-term memory formation and associative learning. I find it attractive to envision the brain as functioning by transitioning back and forth between sparse and distributed coding schemes across different regions, according to the computations needed (Baum et al., 1988; Földiák, 1990, 2002). The design principles emerging from the present study suggest a neural algorithm, or a recipe, of how such transition may be (biologically and algorithmically) accomplished.

REGAINING COMPLEXITY: RE-EXAMINING THE MODEL'S INITIAL ASSUMPTIONS

The model I presented here relies on several rather crude approximations and assumptions (set forth in Section "Model Assumptions"). I now re-examine them in light of experimental data, and wherever they deviate from biological reality, try to assess how the model's predictions are affected. In other words, it's time to make things complicated again.

My first assumption was using discrete time windows, during which PN-spiking is treated as binary (firing or not). The locust olfactory system operates with an internal 20-Hz clock imposed on it (Laurent and Davidowitz, 1994; Laurent and Naraghi, 1994; Perez-Orive et al., 2002, 2004); PNs rarely fire more than once per 50 ms cycle (Perez-Orive et al., 2002, 2004), with odor-evoked spikes confined to the 25 ms rising-phase of the local field-potential oscillation (Laurent and Davidowitz, 1994; Laurent et al., 1996; Wehr and Laurent, 1996). The assumption is thus justified, at least for odor conditions. During baseline the coherence of the PN population is much reduced (as reflected by local-field potentials), yet most PNs retain a 20-Hz oscillatory component, as spike-autocorrelations show (Jortner et al., 2007). On average, the minimal PN-inter-spike-interval during baseline is ∼22 ms; so there are definitely sometimes two spikes per arbitrary 50-ms window, although rarely more than two (Jortner, 2009). This deviation from model assumptions could increase the number of EPSPs summing per time window in a KC during baseline, and reduce the number of inputs needed for threshold-crossing. Having said this, the proportion of spikes with inter-spike-intervals below 50 ms is small, and as PN–KC synapses show no homo-synaptic facilitation on these time scales (Jortner et al., 2007), biases resulting from two spikes per window are expected to be small and linear (at most) with spike number.

A second assumption was i.i.d. spiking across different PNs over the course of the integration time—in other words, that PN activation-patterns are entirely random. This assumption simplified calculations and allowed applying the CLT to the summation of inputs. In reality, however, not all antennal-lobe states are equally probable given that the animal operates in a natural olfactory environment; in fact each individual locust is likely to experience in its lifetime only a miniscule fraction of the enormous number of possible PN activation patterns. Furthermore, during odor presentation PNs are affected by common excitatory input from ORNs and common inhibitory input from local interneurons, making randomness and mutual independence even less likely. In a nutshell, caveats of dependences and correlations between PNs may bias my analyses.

Several points support the model's conclusions despite this potential bias. First, simultaneous recordings show that spikes from different PNs are not correlated over short time scales at baseline (Jortner et al., 2007); this does not establish mutual independence but takes a step in that direction. Second, no direct synaptic connections were ever found between PNs (Jortner and Laurent, in preparation), which eliminates causality from contributing to statistical dependence. Finally, the classical CLT was extended for cases with dependences between variables (Bernstein, 1922, 1927), yielding modern versions of the theorem which hold under various mutual dependences and correlations (e.g., French and Wilson, 1978; Wilson, 1981; Reichert and Schilling, 1985; Pinske et al., 2007). The deviation from i.i.d. firing statistics needs, however, to be borne in mind.

As a third assumption, all PN–KC synaptic connections were treated as equal in strength. Experiments show PN–KC-EPSP amplitudes are distributed narrowly, but they are by no means uniform ($86 \pm 44\,\mu V$, with half of them within 60–110 μV; Jortner et al., 2007). Can ignoring the weights' distribution be justified? All calculations throughout this study were based on summation of rows of the connectivity matrix $\overleftrightarrow{W}$. If the number of summed elements n is large enough, the CLT justifies treating them as uniform (assuming i.i.d. between connections and finite variance, which is reasonable). This holds for "large enough" n, but is the length of a connectivity vector, or the number of EPSPs summating in a target neuron "large enough"? While the CLT strictly applies only when n approaches infinity, it in fact converges to Normality very fast as n starts to increase, then slows down asymptotically (established by the Berry–Esseen theorem: the difference between any CDF with finite variance and the Normal CDF decreases as $1/\sqrt{n}$; Feller, 1972). The number of summated connections in the model is on the order of $N \cdot c$ for connection vectors, and on the order of $N \cdot c \cdot p$ for aggregate inputs—so for large N (800, in our case) the assumption is justified for all but very small values of c or of $c \cdot p$ (with locust parameters, $N \cdot c$ lies within the hundreds, and $N \cdot c \cdot p$ is on the order of tens).

The last point has also been addressed directly with simulations in which sets of EPSPs from the experimental amplitude distribution were randomly drawn and summed (Jortner et al., 2007); for numbers $\geq$50, a Gaussian hypothesis for the sum could no longer be rejected, and differences between the actual sum and its estimate assuming uniformity were minute (Jortner et al., 2007, Supplementary Material). In conclusion, the model results are only minimally biased by the assumption of uniform connection strengths because the numbers are sufficiently high; this must be reexamined, however, if applying this framework to different systems where parameter values may be lower.

Fourth and last, PN–KC connectivity was treated as random. Previous work showed no obvious pattern in the pairs tested positive for connections; in fact, most KCs tested simultaneously with several PNs were found to be connected to about half of them (Jortner et al., 2007). Anatomically, the mushroom body calyx shows no simple patterning (e.g., layered- or

columnar-organization) for either PN axons or KC dendrites (Farivar, 2005). Thus, no data has suggested patterning in the connectivity matrix. This having been said, it is very difficult (in fact, impossible) to establish true randomness in experimental data, and some patterns may have evaded my analysis. Even in such case, however, due to the huge number of combinations $\left(\begin{bmatrix}800\\400\end{bmatrix} \approx 10^{240}\right)$ suggested by the data, any component of randomness in the connectivity matrix would still yield a combinatorial explosion of wiring possibilities—so very dramatic connection biases would be required to alter the conclusions of this study.

RELATED MODELS AND ALTERNATIVE DESIGNS

In this study I took the gross structure of the locust olfactory system as starting point and basis for exploration; I did not explore all possible architectures or parameters, and by no means claim to offer the only solution for constructing specific representations from noisy input. A number of theoretical studies have tackled similar problems using different approaches, and have come up with a variety of designs. Here I briefly survey some of these models and their key properties, comparing and contrasting them with mine.

One example is Kanerva's (1988) Sparse Distributed Memory model. Its central idea—that memories can be represented as binary vectors in a high-dimensional space—stems from a key property of such spaces: points in them tend to differ from each other—and from most of the remaining space—along many dimensions. With this inherent sparseness in mind, a hyper-sphere is drawn around each point of interest (memory), and the memory is activated whenever an input vector falls within the hyper-sphere's boundaries; this grants the model noise-tolerance and flexibility more characteristic of brains than of most computers. As different hyper-spheres may partially overlap, input vectors often activate multiple memories. While performance depends on dimensionality, number of memories stored and activation radii, the model's main results are the feasibility and robustness of sparse-distributed storage and retrieval.

The Sparse Distributed Memory model is fully connected, meaning that when classifying an input, its values along all dimensions are taken into account; zeros as well as ones. This conveys the model its robustness, capacity, and noise tolerance. The threshold (the radius of the hyper-sphere) can be adapted if needed: for example, to ensure that state-space is tiled, or that each output responds with a particular probability. The high-dimensional space is thus filled with many partially overlapping hyper-spheres of the same dimension as the space; each represents one memory and its noise-tolerant envelope.

At another end of the spectrum of connectivity values is Jaeckel's (1989) Selected-Coordinate Design. In this model inputs also reside in a binary, high-dimensional space, yet each output samples just a handful of inputs (10 of 1000; corresponding to connection probability 0.01). For a memory to be activated, all of its sampled inputs (or selected coordinated) must take particular binary values; the rest of the inputs do not matter. Jaeckel's model thus attains its noise tolerance via invariance to most of the input's dimensions: it only takes into account 10 and ignores the rest. The threshold is thus fixed, and is equal to the number of selected coordinates (they all need to be active). In the Selected-Coordinate Design the high-dimensional space is thus inhabited by subspaces of lower dimensionality, each corresponding to a memory. To think in three dimensions, if input space were a cube, memories would be faces (squares) of this cube.

To compare my model with these, it is useful to speak a common language. In the locust, input space has 800 dimensions (one for each PN), so it is also high-dimensional and binary (as I treat each PN as spiking or not within each time window); PN–KC connectivity vectors correspond to points of interest in this space. The interesting properties of my model rely precisely on the inherent sparseness of high-dimensional binary spaces as formulated by Kanerva: PN–KC connectivity vectors populate a space so vast (containing roughly 10^{240} potential points) that the actual 5×10^4 points realized tend to populate it extremely sparsely, each sitting on average very far from all others.

What does the portion of space which KCs respond to look like in my model, and how does their threshold affect it? In Kanerva's model each KC samples all the dimensions and is rather tolerant to errors in any of them (via the threshold); in Jaeckel's model it samples only very few dimensions, but is very strict about perfectly matching these. For comparison, in my model each KC samples half of the dimensions (corresponding to $c = ½$) and is invariant to the rest. This means that in 800-dimensional space, around half of the dimensions—those PNs which the KC is connected to—are treated as spherical, with a threshold, and the others are ignored, thus treated as cubical. The receptive range of a KC will thus be an 800-dimensional hyper-cylinder: spherical along some dimensions and invariant to the others. Adapting the threshold will only affect the spherical dimensions: the larger the radius, the lower the threshold, and thus the more states of the PN population activate the KC. The model suggested here thus combines the high-dimensionality and dense connectivity of Sparse Distributed Memory, the invariance to non-connected inputs from the Selected-Coordinate Design, and elements of noise-tolerance from both.

WHERE ELSE MAY THESE PRINCIPLES APPLY?

Neuronal specificity and sparse coding are widespread phenomena; relevant way beyond olfaction or sensory systems. The design principles discussed here are general in nature, relying on general assumptions and independent of particulars of the system. It is attractive to hypothesize that they may apply in a variety of other interconnected systems. While detailed data on network architecture—especially connection probabilities—is unfortunately still scarce for most biological networks, I point out several candidates which merit comparison.

What happens in other olfactory systems? In *Drosophila*, KCs are concentration invariant and much more specific than PNs (Turner et al., 2008; Honegger et al., 2011). KCs each seem to receive connections from around 10 PNs (corresponding to connectivity of ∼5%), and have high firing thresholds (Turner et al., 2008). Differences in design and coding between locusts and flies may relate to their ecology: fruit flies occupy highly specialized ecological niches whereas locusts are generalist feeders.

KC numbers also differ greatly across these species (50,000 in locust vs. 2500 in *Drosophila*); this could merely reflect size constraints, but may also relate to the extent of odor space the mushroom body needs to tile, or to the resolution required at different regions of the space.

Mammalian pyriform (olfactory) cortex shows similarities with the mushroom body: pyriform pyramidal neurons respond to odors with few spikes locked to respiratory oscillations and have low baseline firing rates (Poo and Isaacson, 2009). Pyriform cortex shows no evidence for spatial organization by odor tuning (Illig and Haberly, 2003; Rennaker et al., 2007; Stettler and Axel, 2009), and axons from individual mitral cells—its input neurons, analogous to insect-PNs—project onto it diffusely, without apparent spatial preference (Friedrich, 2011; Ghosh et al., 2011; Miyamichi et al., 2011; Sosulski et al., 2011). Connection probabilities between mitral cells and pyriform pyramidal cells are unknown; however, as these synapses are strong, coincident input from just a few may suffice to elicit spiking (Franks and Isaacson, 2006). This implies—albeit indirectly—that connection probabilities from second- to third-order neurons in rodents may be lower than in the locust. The level of sparseness of pyriform pyramidal neurons is 3–15%—also considerably lower than in locust KCs (Poo and Isaacson, 2009; Stettler and Axel, 2009; Isaacson, 2010). It remains to be seen how the various network parameters work in concert to yield coding solutions in this system.

One system often treated as a benchmark for decorrelation of representations is the cerebellum. Theoretical work by Marr (1969) and Kanerva (1988) suggests that the transformation from mossy fibers onto cerebellar granule cells is designed to decorrelate input representations and reduces the number of nodes learning would have to act on; operations precisely suited for a neural architecture such as described here. Measurements, however, indicate that convergence ratios of mossy fibers onto granule cells are much lower (Chadderton et al., 2004); the architecture described here may thus not apply to those neurons.

A fascinating candidate for comparison is the mammalian hippocampus. The ability to build meaningful representation from discrete random percepts makes sparse codes attractive for memory formation (Palm, 1980; Baum et al., 1988; Kanerva, 1988, 1993). This is a well-established role of both hippocampus (Scoville and Milner, 1957; Squire, 1992; Tulving and Markowitch, 1998) and mushroom body [reviewed in Heisenberg (1998)], and the analogy between the two has been previously drawn (Strausfeld et al., 1998). Indeed, similarly to KCs, some hippocampal neurons use extremely sparse codes: spiking specifically and reliably in response to complex, high-level stimuli and very rarely at baseline (Kreiman et al., 2000; Barnes et al., 2003; Quian Quiroga et al., 2005), with a majority silent at any given time (Thompson and Best, 1989). Topologically, hippocampus is largely feed-forward (Andersen et al., 1971; O'Reilly and McClelland, 1994; Andersen et al., 2000), and while its cytoarchitecture is extensively studied with classical anatomical techniques (e.g., Amaral and Witter, 1989; Patton and McNaughton, 1995), quantitative functional connectivity-data at single-cell resolution is just emerging (e.g., Brivanlou et al., 2004).

O'Reilly and McClelland (1994) provide in-depth theoretical analysis of hippocampal circuitry. They modeled feed-forward components of the circuit (entorhinal cortex, dentate gyrus, and CA3), exploring the effects of network parameters on pattern-separation and pattern-completion. Testing three values of feed-forward connectivity (equivalent to connection probability ∼0.0001, 0.02, and 0.1), they indeed find that contrary to their intuition, the lowest connectivity value—which they had presumed to outperform the higher ones in pattern-separation—actually performed worse. They found performance similar between the higher values, suggesting a diminishing-returns effect; they did not, however, test higher connectivity-values approaching ∼0.5. It would be tempting to test whether architecture within or among some hippocampal sub-regions (for example CA3–CA1) may follow similar design to the locust PN–KC circuitry, to maximize input separation as a basis for memory formation.

As the design discussed largely relies on random connectivity, it is not inherently suitable for circuits where the input's spatial relations must be retained, such as early visual- or auditory-areas. It may, however, apply well locally within spatially dependent modules, such as cortical columns (Mountcastle, 1997), or in higher processing areas, where representations become object-based and spatially invariant—such as infero-temporal cortex (Gross et al., 1972; Perrett et al., 1982; Fujita et al., 1992; Tanaka, 1996, 2003).

Optimal input-space separation may be useful even when sparse coding is not the goal: the targets' firing threshold determines response probability; if it is low, neurons will respond broadly. For example, connectivity ½ and a low firing threshold can generate distributed representations from sparse ones.

Core mechanisms elucidated here may still apply even with connectivity somewhat removed from the optimum: large enough cell-numbers, intermediate connectivity and some inherent randomness lead to a combinatorial explosion of wiring possibilities. This in turn naturally results in input spaces which are (by virtue of their mere size) extremely sparsely populated. A central message of this study is that to attain efficient input-spread, a suitable source–target connectivity regime is neither very dense, nor very sparse, but rather within an intermediate range.

CONCLUDING REMARKS: ORIGINS OF NEURONAL SPECIFICITY AND THE PARSING OF THE OLFACTORY WORLD

Neuronal specificity and sparse neural coding have continuously attracted attention over several decades of brain research (e.g., Attneave, 1954; Marr, 1969, 1970; Willshaw and Longuet-Higgins, 1970; Barlow, 1972; Palm, 1980; Baum et al., 1988; Kanerva, 1988; Tsodyks and Feigel'Man, 1988; Perez-Vicente and Amit, 1989; Földiák, 1990; Rolls and Tovee, 1995; Vinje and Gallant, 2000; Willmore and Tolhurst, 2001; Simoncelli and Olshausen, 2001; Hahnloser et al., 2002; Laurent, 2002; Perez-Orive et al., 2002; Garcia-Sanchez and Huerta, 2003; DeWeese et al., 2003; Olshausen and Field, 2004; Huerta et al., 2004; Quian Quiroga et al., 2005; Jortner et al., 2007). One reason may be that they highlight a truly fundamental property of the brain: the ability to parse the surroundings and to extract meaning from them. Indeed, it seems that once a network of neurons can—through integration of external sensory inputs and a series of computations—bring single target-cells to respond differentially and reliably to particular objects, combinations of

features or classes of stimuli, a significant part of the way towards performing the brain's tasks has already been made. A set of such "meaningfully responding" cells constitutes the very internal model of the world in the organism's brain—molded to its ecologically dictated requirements and reflecting the world as the animal views it. For example, characterization of a cell ensemble which represents a complex percept, such as an *Apple*, puts a handle on what thinking of an *Apple* is (Barlow, 1972); and strengthening a set of connections between this cell ensemble and another representing the concept of *Cake* creates both an associative link, and a higher, more complex idea. Mechanistic insights into how such representations come into being can open an intimate window onto the brain's subjective world-view and what forms it.

The system I have analyzed here does not yet offer direct access to this level of meaningful representations, but it does highlight the principles on the basis of which they can emerge. The principles along which the olfactory circuitry between the antennal-lobe and mushroom body is designed in the locust are an increase in dimensionality between source- and target-populations, feed-forward connectivity with probability of ½, maximizing separation between representations; and a high and adaptive firing threshold. This leads to specific, reliable, and sparse representations of random olfactory percepts in the mushroom body. Specificity is explained by a high enough threshold, only crossed when the KC encounters an appropriate input vector (from a set of vectors which lie within a particular radius of Hamming distances from an "ideal" central vector); very different from vectors which drive other KCs. Reliability results from the combination of strong convergence of PNs onto KCs (400:1) and cycle-by-cycle adjustment of the KC firing threshold (Papadopoulou et al., 2011).

In the following relay subsets of KCs are sampled by extrinsic β-lobe cells through sparse and strong synapses which are highly plastic (Cassenaer and Laurent, 2007); this can be used to build and learn meaningful representations for β-lobe neurons, constructed from the sparse discrete percepts randomly assigned to KCs (Barlow, 1972; Földiák, 2002). Cells responding to "meaningful" stimuli (for example, plants with high protein-content, or toxic plants) can directly activate motor programs—causing the insect to respond to the stimulus with an appropriate behavior (for example foraging or avoidance, respectively).

The locust olfactory circuitry emerges from this study as general-purpose machinery for information processing: a neural module which receives highly distributed and noisy inputs, spreads them maximally, and creates from them an arbitrarily sparse and selective set of representations—to be used as a substrate for learning, memory formation, categorization/generalization, triggering behavioral programs, and potentially a variety of other computations. These principles are suited to process any input where spatial relations need not be conserved, as they depend only to a limited extent on the nature of the signals to be processed. These mechanisms are therefore potentially of broad applicability and interest; where else they may apply remains to be seen.

ACKNOWLEDGMENTS

I am grateful to Gilles Laurent, Idan Segev, Markus Meister, Bill Bialek, and Gilad Jacobson for fruitful discussions during the course of this work. I thank the reviewers for their comments and suggestions, and especially Pentti Kanerva for his in-depth, critical and highly constructive review. This work was mostly done at the Hebrew University in Jerusalem and was partially funded by the Horowitz Foundation.

REFERENCES

Amaral, D. G., and Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience* 31, 571–591.

Andersen, P., Bliss, T. V. P., and Skrede, K. (1971). Lamellar organization of hippocampal excitatory pathways. *Exp. Brain Res.* 13, 222–238.

Andersen, P., Soleng, A., and Raastad, M. (2000). The hippocampal lamella hypothesis revisited. *Brain Res.* 886, 165–171.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychol. Rev.* 61, 183–193.

Attwell, D., and Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metab.* 21, 1133–1145.

Barlow, H. B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1, 371–394.

Barlow, H. B. (1985). "Cerebral cortex as a model builder," in *Models of the Visual Cortex*, eds D. Rose and V. G. Dobson (Chichester: Wiley), 37–46.

Barnes, C. A., Skaggs, W. E., McNaughton, B. L., Haworth, M. L., Permenter, M., Archibeque, M., et al. (2003). "Chronic recording of neuronal populations in the temporal lobe of awake young adult and geriatric primates," in *Program no. 518.8. 2003 Neuroscience Meeting Planner* (New Orleans, LA: Society for Neuroscience).

Baum, E. B., Moody, J., and Wilczek, F. (1988). Internal representations for associative memory. *Biol. Cybern.* 59, 217–228.

Behmer, S. T., Belt, C. E., and Shapiro, M. S. (2005). Variable rewards and discrimination ability in an insect herbivore: what and how does a hungry locust learn? *J. Exp. Biol.* 208, 3463–3473.

Behmer, S. T., Elias, D. O., and Bernays, E. A. (1999). Post-ingestive feedbacks and associative learning regulate the intake of unsuitable sterols in a generalist grasshopper. *J. Exp. Biol.* 202, 739–748.

Bernays, E. A., Bright, K., Howard, J. J., Raubenheimer, D., and Champagne, D. (1992). Variety is the spice of life: frequent switching between foods in the polyphagous grasshopper, *Taeniopoda eques*. *Anim. Behav.* 44, 721–731.

Bernays, E. A., and Chapman, R. F. (1994). *Host-Plant Selection by Phytophagous Insects*. New York, NY: Chapman and Hall.

Bernstein, S. (1922). Sur le théorème limit du calcul des probabilités. *Math. Ann.* 85, 237–241.

Bernstein, S. (1927). Sur l'extension du théorème du calcul des probabilités aux sommes de quantités dependantes. *Math. Ann.* 97, 1–59.

Brivanlou, I. H., Dantzker, J. L. M., Stevens, C. F., and Callaway, E. M. (2004). Topographic specificity of functional connections from hippocampal CA3 to CA1. *Proc. Natl. Acad. Sci. U.S.A.* 101, 2560–2565.

Cassenaer, S., and Laurent, G. (2007). Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* 448, 709–714.

Chadderton, P., Margrie, T. W., and Häusser, M. (2004). Integration of quanta in cerebellar granule cells during sensory processing. *Nature* 428, 856–860.

Churchland, P. S., and Sejnowski, T. J. (1990). Neural representation and neural computation. *Philos. Perspect.* 4, 343–382.

Churchland, P. S., and Sejnowski, T. J. (1992). *The Computational Brain*. Cambridge, MA: MIT Press.

Cooper-Driver, G., Swain, T., Bernays, E. A., and Finch, S. (1977). Seasonal variation in secondary plant compounds in relation to palatability of *Pteridium aquilinum*. *Biochem. Syst. Ecol.* 5, 177–183.

Cottee, P., Bernays, E. A., and Mordue, A. J. (1988). Comparisons of deterrency and toxicity of selected secondary plant compounds to an oligophagous and a polyphagous acridid. *Entomol. Exp. Appl.* 46, 241–247.

deCharms, R. C., and Zador, A. (2000). Neural representation and the cortical code. *Annu. Rev. Neurosci.* 23, 613–647.

DeWeese, M. R., Wehr, M., and Zador, A. M. (2003). Binary spiking in auditory cortex. *J. Neurosci.* 23, 7940–7949.

Dukas, R., and Bernays, E. A. (2000). Learning improves growth in the grasshopper, *Schistocerca americana*. *Proc. Natl. Acad. Sci. U.S.A.* 97, 2637–2640.

Farivar, S. S. (2005). *Cytoarchitecture of the Locust Olfactory System*. Ph.D. Thesis, California Institute of Technology, Pasadena, CA.

Feller, W. (1972). *An Introduction to Probability Theory and its Applications, Vol. II, 2nd Edn.* New York, NY: John Wiley and Sons.

Finelli, L. A., Haney, S., Bazhenov, M., Stopfer, M., and Sejnowski, T. J. (2008). Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput. Biol.* 4:e1000062. doi: 10.1371/journal.pcbi.1000062

Földiák, P. (1990). Forming sparse representations by local anti-Hebbian learning. *Biol. Cybern.* 64, 165–170.

Földiák, P. (2002). "Sparse coding in the primate cortex," in *The Handbook of Brain Theory and Neural Networks, 2nd Edn,* ed M. A. Arbib (Cambridge, MA: MIT Press), 1064–1067.

Franks, K. M., and Isaacson, J. S. (2006). Strong single-fiber sensory inputs to olfactory cortex: implications for olfactory coding. *Neuron* 49, 357–363.

French, S., and Wilson, K. (1978). Treatment of a negative intensive observation. *Acta Cryst. Sect. A* 34, 517–525.

Friedrich, R. W. (2011). Olfactory neuroscience: beyond the bulb. *Curr. Biol.* 21, 438–440.

Fujita, I., Tanaka, K., Ito, M., and Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360, 343–346.

Garcia-Sanchez, M., and Huerta, R. (2003). Design parameters of the fan-out phase of sensory systems. *J. Comput. Neurosci.* 15, 5–17.

Ghosh, S., Larson, S. D., Hefzi, H., Marnoy, Z., Cutforth, T., Dokka, K., et al. (2011). Sensory maps in the olfactory cortex defined by long-range viral tracing of single neurons. *Nature* 472, 217–220.

Gross, C. G., Rocha-Miranda, C., and Bender, D. (1972). Visual properties of neurons in the inferotemporal cortex of the macaque. *J. Neurophysiol.* 35, 96–111.

Hahnloser, R. H., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419, 65–70.

Hammer, M., and Menzel, R. (1995). Learning and memory in honeybees. *J. Neurosci.* 15, 1617–1630.

Hammer, M., and Menzel, R. (1998). Multiple sites of associative odor learning as revealed by local brain microinjections of octopamine in honeybees. *Learn. Mem.* 5, 146–156.

Heisenberg, M. (1998). What do the mushroom bodies do for the insect brain? an introduction. *Learn. Mem.* 5, 1–10.

Honegger, K. S., Campbell, R. A., and Turner, G. C. (2011). Cellular-resolution population imaging reveals robust sparse coding in the *Drosophila* mushroom body. *J. Neurosci.* 31, 11772–11785.

Hopkins, T. L., and Young, H. (1990). Attraction of the grasshopper, *Melanoplus sanguinipes*, to host plant odors and volatile components. *Entomol. Exp. Appl.* 56, 249–258.

Huerta, R., Nowotny, T., Garcia-Sanchez, M., Abarbanel, H. D., and Rabinovich, M. I. (2004). Learning classification in the olfactory system of insects. *Neural Comput.* 16, 1601–1640.

Illig, K. R., and Haberly, L. B. (2003). Odor-evoked activity is spatially distributed in piriform cortex. *J. Comp. Neurol.* 457, 361–373.

Isaacson, J. S. (2010). Odor representations in mammalian cortical circuits. *Curr. Opin. Neurobiol.* 20, 328–331.

Jaeckel, L. A. (1989). "An alternative design for a sparse distributed memory," in *RIACS Technical Report TR 89.28* (Mountain View, CA: Research Institute for Advanced Computer Science, NASA Ames Research Center).

Jortner, R. A. (2009). *Linking Network Architecture and Neural Coding in the Olfactory System of the Locust*. Ph.D. Thesis, The Hebrew University, Jerusalem.

Jortner, R. A., Farivar, S. S., and Laurent, G. (2007). A simple connectivity scheme for sparse coding in an olfactory system. *J. Neurosci.* 27, 1659–1669.

Kadohisa, M., and Wilson, D. A. (2006). Olfactory cortical adaptation facilitates detection of odors against background. *J. Neurophysiol.* 95, 1888–1896.

Kanerva, P. (1988). *Sparse Distributed Memory*. Cambridge, MA: MIT Press.

Kanerva, P. (1993). "Sparse distributed memory and related models," in *Associative Neural Memories: Theory and Implementation*, ed M. H. Hassoun (New York, NY: Oxford University Press), 50–76.

Kreiman, G., Koch, C., and Fried, I. (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat. Neurosci.* 3, 946–953.

Laurent, G. (1996). Dynamical representation of odors by oscillating and evolving neural assemblies. *Trends Neurosci.* 19, 489–496.

Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nat. Rev. Neurosci.* 3, 884–895.

Laurent, G., and Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science* 265, 1872–1875.

Laurent, G., and Naraghi, M. (1994). Odorant-induced oscillations in the mushroom bodies of the locust. *J. Neurosci.* 14, 2993–3004.

Laurent, G., Wehr, M., and Davidowitz, H. (1996). Temporal representation of odors in an olfactory network. *J. Neurosci.* 16, 3837–3847.

Leitch, B., and Laurent, G. (1996). GABAergic synapses in the antennal lobe and mushroom body of the locust olfactory system. *J. Comp. Neurol.* 372, 487–514.

Levy, W. B., and Baxter, R. A. (1996). Energy efficient neural codes. *Neural Comput.* 8, 531–543.

Li, Y. S., and Strausfeld, N. J. (1997). Morphology and sensory modality of mushroom body extrinsic neurons in the brain of the cockroach, *Periplaneta americana*. *J. Comp. Neurol.* 387, 631–650.

MacLeod, K., Bäcker, A., and Laurent, G. (1998). Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 395, 693–698.

Marr, D. (1969). A theory of cerebellar cortex. *J. Physiol.* 202, 437–470.

Marr, D. (1970). A theory for cerebral neocortex. *Proc. R. Soc. Lond. B Biol. Sci.* 176, 161–234.

Marr, D. (1971). Simple memory: a theory for archicortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 262, 23–81.

Mazor, O., and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48, 661–673.

Miyamichi, K., Amat, F., Moussavi, F., Wang, C., Wickersham, I., Wall, N. R., et al. (2011). Cortical representations of olfactory input by trans-synaptic tracing. *Nature* 472, 191–196.

Mountcastle, V. (1997). The columnar organization of the neocortex. *Brain* 120, 701–722.

Nowotny, T. (2009). "Sloppy engineering and the olfactory system of insects," in *Biologically Inspired Signal Processing for Chemical Sensing, Studies in Computational Intelligence, Vol. 188,* eds S. Marco and A. Gutiérrez (Berlin, Heidelberg: Springer), 3–32.

Olshausen, B. A., and Field, D. J. (2004). Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* 14, 481–487.

O'Reilly, R. C., and McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage and recall: avoiding a tradeoff. *Hippocampus* 4, 661–682.

Palm, G. (1980). On associative memory. *Biol. Cybern.* 59, 217–228.

Papadopoulou, M., Cassenaer, S., Nowotny, T., and Laurent, G. (2011). Normalization for sparse encoding of odors by a wide-field interneuron. *Science* 332, 721–725.

Patton, P. E., and McNaughton, B. (1995). Connection matrix of the hippocampal formation: I. the dentate gyrus. *Hippocampus* 5, 245–286.

Perez-Orive, J., Bazhenov, M., and Laurent, G. (2004). Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *J. Neurosci.* 24, 6037–6047.

Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science* 297, 359–365.

Perez-Vicente, C. J., and Amit, D. J. (1989). Optimised network for sparsely coded patterns. *J. Phys. A* 22, 559–569.

Perrett, D. I., Rolls, E. T., and Caan, W. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Exp. Brain Res.* 47, 329–342.

Pinske, J., Shen, L., and Slade, M. (2007). A central limit theorem for endogenous locations and complex spatial interactions. *J. Econometr.* 140, 215–225.

Poo, C., and Isaacson, J. S. (2009). Odor representations in olfactory cortex: "sparse" coding, global inhibition, and oscillations. *Neuron* 62, 850–861.

Quian Quiroga, R., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single-neurons in the human brain. *Nature* 435, 1102–1107.

Raman, B., and Stopfer, M. (2010). Analysis of trial-by-trial variability in stimulus-evoked neural activity. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* 2010, 4320–4322.

Raman, B., Stopfer, M., and Semancik, S. (2011). Mimicking biological design and computing principles in artificial olfaction. *ACS Chem. Neurosci.* 2, 487–499.

Reichert, P., and Schilling, R. (1985). A local limit theorem for strongly dependent random variables and its application to a chaotic configuration of atoms. *J. Math. Phys.* 26, 1165–1172.

Rennaker, R. L., Chen, C. F., Ruyle, A. M., Sloan, A. M., and Wilson, D. A. (2007). Spatial and temporal distribution of odorant-evoked activity in the piriform cortex. *J. Neurosci.* 27, 1534–1542.

Rolls, E. T., and Tovee, M. J. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73, 713–726.

Schwaerzel, M., Monastirioti, M., Scholz, H., Friggi-Grelin, F., Birman, S., and Heisenberg, M. (2003). Dopamine and octopamine differentiate between aversive and appetitive olfactory memories in *Drosophila*. *J. Neurosci.* 23, 10495–10502.

Scoville, W. B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* 20, 11–21.

Simoncelli, E. P., and Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci.* 24, 1193–1216.

Sivan, E., and Kopell, N. (2004). Mechanism and circuitry for clustering and fine discrimination of odors in insects. *Proc. Natl. Acad. Sci. U.S.A.* 101, 17861–17866.

Sosulski, D. L., Bloom, M. L., Cutforth, T., Axel, R., and Datta, S. R. (2011). Distinct representations of olfactory information in different cortical centres. *Nature* 472, 213–216.

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys and humans. *Psychol. Rev.* 99, 195–231.

Stettler, D. D., and Axel, R. (2009). Representations of odor in the piriform cortex. *Neuron* 63, 854–864.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.

Strausfeld, N. J., Hansen, L., Li, Y., Gomez, R. S., and Ito, K. (1998). Evolution, discovery, and interpretations of arthropod mushroom bodies. *Learn. Mem.* 5, 11–37.

Stryker, M. (1992). Elements of visual perception. *Nature* 360, 301–302.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annu. Rev. Neurosci.* 19, 109–139.

Tanaka, K. (2003). Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. *Cereb. Cortex* 13, 90–99.

Theunissen, F. E. (2003). From synchrony to sparseness. *Trends Neurosci.* 26, 61–64.

Thompson, L. T., and Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *J. Neurosci.* 9, 2382–2390.

Tsodyks, M. V., and Feigel'Man, M. V. (1988). The enhanced storage capacity in neural networks with low activity level. *Neurophys. Lett.* 6, 101–105.

Tulving, E., and Markowitch, H. J. (1998). Episodic and declarative memory: role of the hippocampus. *Hippocampus* 8, 198–204.

Turner, G. C., Bazhenov, M., and Laurent, G. (2008). Olfactory representations by *Drosophila* mushroom body neurons. *J. Neurophysiol.* 99, 734–746.

Unoki, S., Matsumoto, Y., and Mizunami, M. (2005). Participation of octopaminergic reward system and dopaminergic punishment system in insect olfactory learning revealed by pharmacological study. *Eur. J. Neurosci.* 22, 1409–1416.

Vinje, W. E., and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287, 1273–1276.

von der Malsburg, C. (1986). "Am I thinking assemblies?" in *Brain Theory,* eds G. Palm and A. Aertsen (Berlin, Heidelberg, New York, NY: Springer), 161–176.

von der Malsburg, C. (1990). "A neural architecture for the representation of scenes," in *Brain Organization and Memory: Cells, Systems and Circuits*, eds J. L. McGaugh, N. M. Weinberger, and G. Lynch (New York, NY: Oxford University Press), 356–372.

Wehr, M., and Laurent, G. (1996). Odor encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384, 162–166.

Willmore, B., and Tolhurst, D. J. (2001). Characterizing the sparseness of neural codes. *Network* 12, 255–270.

Willshaw, D. J., and Longuet-Higgins, H. C. (1970). Associative memory models. *Mach. Intell.* 5, 351–359.

Wilson, A. J. C. (1981). Can intensity statistics accommodate stereochemistry? *Acta Cryst. Sect. A* 37, 808–810.

Conflict of Interest Statement: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 14 October 2011; paper pending published: 07 December 2011; accepted: 20 November 2012; published online: 03 January 2013.

*Citation: Jortner RA (2013) Network architecture underlying maximal separation of neuronal representations. Front. Neuroeng. **5**:19. doi: 10.3389/fneng.2012.00019*

APPENDIX (JORTNER, 2012)

A1. HAMMING DISTANCE BETWEEN PN–KC CONNECTIVITY VECTORS—FULL DERIVATION

Let $\overrightarrow{U}$, $\overrightarrow{V}$ be two arbitrary connectivity vectors, or rows of the connectivity matrix $\overleftrightarrow{W}$ (where 1 denotes connection, with probability c, and 0 none, with probability $1-c$) between the populations $\overrightarrow{S}$ and $\overrightarrow{K}$. The Hamming distance, $H(\overrightarrow{U}, \overrightarrow{V})$, counts the number of bits different across the two vectors. I derive $\left\langle H\left(\overrightarrow{U}, \overrightarrow{V}\right)\right\rangle_{U,V}$, the average Hamming distance between $\overrightarrow{U}$, $\overrightarrow{V}$ over all their possible values:

$$\left\langle H\left(\overrightarrow{U}, \overrightarrow{V}\right)\right\rangle_{U,V} = \left\langle \sum_{i=1}^{N} (U_i - V_i)^2 \right\rangle_{U,V}$$

$$= \left\langle \sum_{i=1}^{N} \left(U_i^2 - 2U_iV_i + V_i^2\right) \right\rangle_{U,V}$$

as the average of a sum is the sum of the averages,

$$= \left\langle \sum_{i=1}^{N} U_i^2 \right\rangle_U - \left\langle \sum_{i=1}^{N} 2U_iV_i \right\rangle_{U,V} + \left\langle \sum_{i=1}^{N} V_i^2 \right\rangle_V$$

$$= \sum_{i=1}^{N} \left\langle U_i^2 \right\rangle_U - 2\sum_{i=1}^{N} \langle U_iV_i \rangle_{U,V} + \sum_{i=1}^{N} \left\langle V_i^2 \right\rangle_V$$

Finishing the calculation now requires the expected values of the expressions U_i^2, V_i^2, U_iV_i. The following table gives all possible values of U_i, U_i^2 and their respective probabilities:

U_i	U_i^2	Probability
1	1	c
0	0	$(1-c)$

So the expected value of $\left\langle U_i^2 \right\rangle_U = 1 \cdot c + 0 \cdot (1-c) = c$; the same holds for $\left\langle V_i^2 \right\rangle_V$.

Here are all possible values of U_i, V_i, U_iV_i and their respective probabilities:

U_i	V_i	U_iV_i	Probability
1	1	1	c^2
1	0	0	$c \cdot (1-c)$
0	1	0	$(1-c) \cdot c$
0	0	0	$(1-c)^2$

So the expected value is

$$\langle V_iU_i \rangle_{U,V} = 1 \cdot c^2 + 0 \cdot \left(c \cdot (1-c) + (1-c) \cdot c + (1-c)^2\right) = c^2$$

Finishing the calculation:

$$\left\langle H(\vec{U}, \vec{V}) \right\rangle_{U,V} = N \cdot c - 2N \cdot c^2 + N \cdot c = 2N \cdot c \cdot (1-c)$$

A2. VARIANCE OF INPUT TO A KC (Λ)—FULL DERIVATION

I substitute the mean input to a KC by Ψ, which was already calculated (Section "Model Results II: Neuronal Activity and Properties of Input to KCs"):

$$\Lambda \equiv \langle \mathrm{var}(k_i) \rangle_i = \left\langle \left\langle \left(k_i - \langle k_i \rangle_{\vec{S}}\right)^2 \right\rangle_{\vec{S}} \right\rangle_i = \left\langle \left\langle (k_i - \Psi)^2 \right\rangle_{\vec{S}} \right\rangle_i$$

$$= \left\langle \left\langle \left(\sum_{j=1}^{N} S_jW_{ij} - \Psi \right)^2 \right\rangle_{\vec{S}} \right\rangle_i$$

$$= \left\langle \left\langle \left(\sum_{j=1}^{N} S_jW_{ij} \right)^2 - 2\Psi \sum_{j=1}^{N} S_jW_{ij} + \Psi^2 \right\rangle_{\vec{S}} \right\rangle_i$$

$$= \left\langle \left\langle \sum_{j=1}^{N}\sum_{k=1}^{N} S_jS_kW_{ij}W_{ik} - 2\Psi \sum_{j=1}^{N} S_jW_{ij} + \Psi^2 \right\rangle_{\vec{S}} \right\rangle_i$$

$$= \left\langle \sum_{j=1}^{N}\sum_{k=1}^{N} \langle S_jS_k \rangle_{\vec{S}} W_{ij}W_{ik} - 2\Psi \sum_{j=1}^{N} \langle S_j \rangle_{\vec{S}} W_{ij} + \Psi^2 \right\rangle_i$$

Next, I'll separate the first-term into its non-diagonal ($i \neq j$) and diagonal ($i = j$) components and treat them each separately:

$$\Lambda = \left\langle \sum_{j=1}^{N}\sum_{k=1,\, j\neq k}^{N} \langle S_jS_k \rangle_{\vec{S}} W_{ij}W_{ik} + \sum_{j=1}^{N} \left\langle S_j^2 \right\rangle_{\vec{S}} W_{ij}^2 - \ldots \right.$$

$$\left. - 2\Psi \sum_{j=1}^{N} \langle S_j \rangle_{\vec{S}} W_{ij} + \Psi^2 \right\rangle_i$$

To calculate the expected values of the non-diagonal terms S_i, S_j, S_iS_j, the table below provides all possible values and their respective probabilities:

S_i	S_j	S_iS_j	Probablity
1	1	1	p^2
1	0	0	$p \cdot (1-p)$
0	1	0	$(1-p) \cdot p$
0	0	0	$(1-p)^2$

So the expected value is

$$\langle S_iS_j \rangle_{\vec{S}, i\neq j} = 1 \cdot p^2 + 0 \cdot \left(p \cdot (1-p) + (1-p) \cdot p + (1-p)^2\right) = p^2$$

For the diagonal terms: S_j, S_j^2 and their respective probabilities:

S_j	S_j^2	Probability
1	1	p
0	0	$(1-p)$

The expected value of $\langle S_j^2\rangle_{\vec{S}} = 1\cdot p + 0\cdot(1-p) = p$, the same holds for $\langle S_j\rangle_{\vec{S}} = 1\cdot p + 0\cdot(1-p) = p$

Continuing the calculation:

$$\Lambda = p^2\cdot\sum_{j=1}^{N}\sum_{k=1,\ j\neq k}^{N}\langle W_{ij}W_{ik}\rangle_i + p\cdot\sum_{j=1}^{N}\left\langle W_{ij}^2\right\rangle_i - \ldots$$
$$-2\Psi\cdot p\cdot\sum_{j=1}^{N}\langle W_{ij}\rangle_i + \Psi^2$$

All possible values for $\langle W_{ij}W_{ik}\rangle_{i,j\neq k}$, $\langle W_{ij}^2\rangle_{i,j}$, $\langle W_{ij}\rangle_{i,j}$ and their respective probabilities:

W_{ij}	W_{ik}	$W_{ij}W_{ik}$	Probability
1	1	1	c^2
1	0	0	$c\cdot(1-c)$
0	1	0	$(1-c)\cdot c$
0	0	0	$(1-c)^2$

So the expected value is

$$\langle W_{ij}W_{ik}\rangle_{i,j\neq k} = 1\cdot c^2 + 0\cdot(c\cdot(1-c) + (1-c)\cdot c + (1-c)^2) = c^2$$

W_{ij}	W_{ij}^2	Probability
1	1	c
0	0	$(1-c)$

the expected value of $\left\langle W_{ij}^2\right\rangle_{i,j} = 1\cdot c + 0\cdot(1-c) = c$

and the same holds for $\langle W_{ij}\rangle_{i,j} = 1\cdot c + 0\cdot(1-c) = c$

There are exactly (N^2-N) non-diagonal terms, and N diagonal terms, so

$$\Lambda = p^2\cdot\left(N^2-N\right)\cdot c^2 + p\cdot N\cdot c - 2\Psi\cdot p\cdot N\cdot c + \Psi^2$$

I will now substitute back $\Psi = N\cdot p\cdot c$ (as shown in Section "Model Results II: Neuronal Activity and Properties of Input to KCs"), and get:

$$\Lambda = \Psi^2 - p^2\cdot N\cdot c^2 + p\cdot N\cdot c - 2\Psi^2 + \Psi^2$$
$$= N\cdot p\cdot c\cdot(1-p\cdot c) = N\cdot p\cdot c\cdot(1-p\cdot c)$$

A3. COVARIANCE OF THE INPUTS TO TWO KCs—FULL DERIVATION

To calculate the covariance between inputs (or between sub-threshold membrane potentials) of two KCs, I substitute the mean input to a KC by Ψ, which was already calculated (Section "Model Results II: Neuronal Activity and Properties of Input to KCs").

$$\langle\mathrm{cov}(k_r,k_t)\rangle_{r\neq t} = \left\langle\left(k_r - \langle k_r\rangle_{\vec{S}}\right)\left(k_t - \langle k_t\rangle_{\vec{S}}\right)\right\rangle_{\vec{S}}\Big\rangle_{r\neq t}$$
$$= \left\langle\left\langle\left(\sum_{i=1}^{N}S_iW_{ri} - \Psi\right)\left(\sum_{j=1}^{N}S_jW_{tj} - \Psi\right)\right\rangle_{\vec{S}}\right\rangle_{r\neq t}$$
$$= \left\langle\left\langle\sum_{i=1}^{N}\sum_{j=1}^{N}S_iS_jW_{ri}W_{tj} - \Psi\sum_{i=1}^{N}S_iW_{ri} - \ldots\right.\right.$$
$$\left.\left.-\Psi\sum_{j=1}^{N}S_jW_{tj} + \Psi^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t}$$
$$= \left\langle\sum_{i=1}^{N}\sum_{j=1}^{N}\langle S_iS_j\rangle_{\vec{S}}\cdot W_{ri}W_{tj} - \ldots\right.$$
$$\left.-\Psi\sum_{i=1}^{N}\langle S_i\rangle_{\vec{S}}\cdot W_{ri} - \Psi\sum_{j=1}^{N}\langle S_j\rangle_{\vec{S}}\cdot W_{tj} + \Psi^2\right\rangle_{r\neq t}$$

separating the first-term into non-diagonal and diagonal components:

$$\left\langle\sum_{i=1}^{N}\sum_{j=1,i\neq j}^{N}\langle S_iS_j\rangle_{\vec{S}}\cdot W_{ri}W_{tj} + \sum_{i=1}^{N}\langle S_i^2\rangle_{\vec{S}}\cdot W_{ri}W_{ti} - \ldots\right.$$
$$\left.-\Psi\sum_{i=1}^{N}\langle S_i\rangle_{\vec{S}}\cdot W_{ri} - \Psi\sum_{j=1}^{N}\langle S_j\rangle_{\vec{S}}\cdot W_{tj} + \Psi^2\right\rangle_{r\neq t}$$

I calculate the expected values of the terms S_iS_j, S_i^2, S_i exactly as in Appendix A2:

$$\left\langle p^2\cdot\sum_{i=1}^{N}\sum_{j=1,\ i\neq j}^{N}W_{ri}W_{tj} + p\cdot\sum_{i=1}^{N}W_{ri}W_{ti} - p\cdot\Psi\sum_{i=1}^{N}W_{ri} - \ldots\right.$$
$$\left.-p\cdot\Psi\sum_{j=1}^{N}W_{tj} + \Psi^2\right\rangle_{r\neq t}$$
$$= p^2\cdot\sum_{i=1}^{N}\sum_{j=1,i\neq j}^{N}\langle W_{ri}W_{tj}\rangle_{r\neq t} + p\cdot\sum_{i=1}^{N}\langle W_{ri}W_{ti}\rangle_{r\neq t} - \ldots$$
$$-p\cdot\Psi\sum_{i=1}^{N}\langle W_{ri}\rangle_{r\neq t} - p\cdot\Psi\sum_{j=1}^{N}\langle W_{tj}\rangle_{r\neq t} + \Psi^2$$

and because of the condition $r\neq t$, the expected value of the term $\langle W_{ri}W_{tj}\rangle_{r\neq t,i\neq j}$ is identical to that of $\langle W_{ri}W_{ti}\rangle_{r\neq t}$, both equal to c^2. The rest of the terms are calculated exactly as in A2, so:

$$p^2\cdot(N^2-N)\cdot c^2 + p\cdot N\cdot c^2 - p\cdot\Psi\cdot N\cdot c - p\cdot\Psi\cdot N\cdot c + \Psi^2$$
$$= \Psi^2 - p^2\cdot N\cdot c^2 + p\cdot N\cdot c^2 - \Psi^2 - \Psi^2 + \Psi^2$$
$$= N\cdot c^2\cdot p\cdot(1-p)$$

A4. DIFFERENCE BETWEEN TWO KC INPUTS (OR SUB-THRESHOLD MEMBRANE POTENTIALS)

Here I follow the exact same lines of reasoning as in Appendix A1–A3. First, I express the difference between KC inputs, second, I split it into its non-diagonal and diagonal components, and third, I average each class of terms separately:

$$
\begin{aligned}
D(k_r,k_t) \equiv \left\langle\left\langle (k_r-k_t)^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t} &= \left\langle\left\langle\left(\sum_{i=1}^{N} S_i W_{ri} - \sum_{j=1}^{N} S_j W_{tj}\right)^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\
&= \left\langle\left\langle\left(\sum_{i=1}^{N} S_i W_{ri} - \sum_{j=1}^{N} S_j W_{tj}\right)\cdot\left(\sum_{m=1}^{N} S_m W_{rm} - \sum_{n=1}^{N} S_n W_{tn}\right)\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\
&= \left\langle\left\langle\sum_{i=1}^{N}\sum_{m=1}^{N} S_i S_m W_{ri} W_{rm}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} - \left\langle\left\langle\sum_{i=1}^{N}\sum_{n=1}^{N} S_i S_n W_{ri} W_{tn}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} - \ldots \\
&\quad - \left\langle\left\langle\sum_{j=1}^{N}\sum_{m=1}^{N} S_j S_m W_{tj} W_{rm}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} + \left\langle\left\langle\sum_{j=1}^{N}\sum_{n=1}^{N} S_j S_n W_{tj} W_{tn}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\
&= \left\langle\left\langle\sum_{i=1}^{N}\sum_{m=1,\, i\neq m}^{N} S_i S_m W_{ri} W_{rm}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} + \left\langle\left\langle\sum_{i=1}^{N} S_i^2\cdot W_{ri}^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t} - \ldots \\
&\quad - \left\langle\left\langle\sum_{i=1}^{N}\sum_{n=1,\, i\neq n}^{N} S_i S_n W_{ri} W_{tn}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} - \left\langle\left\langle\sum_{i=1}^{N} S_i^2 W_{ri} W_{ti}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} - \ldots \\
&\quad - \left\langle\left\langle\sum_{j=1}^{N}\sum_{m=1,\, j\neq m}^{N} S_j S_m W_{tj} W_{rm}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} - \left\langle\left\langle\sum_{j=1}^{N} S_j^2 W_{tj} W_{rj}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} + \ldots \\
&\quad + \left\langle\left\langle\sum_{j=1}^{N}\sum_{n=1,\, j\neq n}^{N} S_j S_n W_{tj} W_{tn}\right\rangle_{\vec{S}}\right\rangle_{r\neq t} + \left\langle\left\langle\sum_{j=1}^{N} S_j^2 W_{tj}^2\right\rangle_{\vec{S}}\right\rangle_{r\neq t} \\
&= \sum_{i=1}^{N}\sum_{m=1,\, i\neq m}^{N} \langle S_i S_m\rangle_{\vec{S}}\cdot\langle W_{ri} W_{rm}\rangle_{r\neq t} + \sum_{i=1}^{N}\left\langle S_i^2\right\rangle_{\vec{S}}\cdot\left\langle W_{ri}^2\right\rangle_{r\neq t} - \ldots \\
&\quad - \sum_{i=1}^{N}\sum_{n=1,\, i\neq n}^{N} \langle S_i S_n\rangle_{\vec{S}}\cdot\langle W_{ri} W_{tn}\rangle_{r\neq t} - \sum_{i=1}^{N}\left\langle S_i^2\right\rangle_{\vec{S}}\cdot\langle W_{ri} W_{ti}\rangle_{r\neq t} - \ldots \\
&\quad - \sum_{j=1}^{N}\sum_{m=1,\, j\neq m}^{N} \left\langle S_j S_m\right\rangle_{\vec{S}}\cdot\left\langle W_{tj} W_{rm}\right\rangle_{r\neq t} - \sum_{j=1}^{N}\left\langle S_j^2\right\rangle_{\vec{S}}\cdot\left\langle W_{tj} W_{rj}\right\rangle_{r\neq t} + \ldots \\
&\quad + \sum_{j=1}^{N}\sum_{n=1,\, j\neq n}^{N} \left\langle S_j S_n\right\rangle_{\vec{S}}\cdot\left\langle W_{tj} W_{tn}\right\rangle_{r\neq t} + \sum_{j=1}^{N}\left\langle S_j^2\right\rangle_{\vec{S}}\cdot\left\langle W_{tj}^2\right\rangle_{r\neq t} \\
&= N(N-1)\cdot p^2\cdot c^2 + N\cdot p\cdot c - N(N-1)\cdot p^2\cdot c^2 - N\cdot p\cdot c^2 - \ldots \\
&\quad - N(N-1)\cdot p^2\cdot c^2 - N\cdot p\cdot c^2 + N(N-1)\cdot p^2\cdot c^2 + N\cdot p\cdot c \\
&= 2N\cdot p\cdot c\cdot(1-c)
\end{aligned}
$$

A5. FIRING THRESHOLD WHEN CONNECTIVITY IS 1/2

A threshold designed to be crossed for a particular fraction of states $1-Q(z)$, where Q is the CDF of the standard Normal distribution, should be equal to the mean (Ψ) plus the appropriate times the standard deviation ($\sqrt{\Lambda}$).

Assuming $c = ½$, we get:

$$
f(z) = \Psi + z\sqrt{\Lambda} = N\cdot p\cdot c + z\sqrt{N\cdot p\cdot c\cdot(1-p\cdot c)}
$$

$$
c = \frac{1}{2} \Rightarrow f(z) = \frac{N\cdot p}{2} + z\sqrt{\frac{N\cdot p}{2}\left(1-\frac{p}{2}\right)} = \frac{N\cdot p + z\sqrt{N\cdot p\cdot\left(2-p\right)}}{2}
$$

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 02 February 2012
doi: 10.3389/fneng.2012.00001

Characterization of a clinical olfactory test with an artificial nose

David J. Yáñez[1], Adolfo Toledano[2], Eduardo Serrano[3], Ana M. Martín de Rosales[4], Francisco B. Rodríguez[3] and Pablo Varona[3]*

[1] *Deutecno Noses S.L., Madrid, Spain*
[2] *Department of Otorhinolaryngology, Fundación Hospital Alcorcón, Madrid, Spain*
[3] *Grupo de Neurocomputación Biológica, Departamento de Ingeniería Informática, Escuela Politécnica Superior, Universidad Autónoma de Madrid, Madrid, Spain*
[4] *Pharmacy Unit, Fundación Hospital Alcorcón, Madrid, Spain*

Edited by:
Ramón Huerta, University of California, USA

Reviewed by:
Alberto Ferrus, Cajal Institute (CSIC), Spain
Roberto F. Galán, Case Western Reserve University, USA

***Correspondence:**
Pablo Varona, Grupo de Neurocomputación Biológica, Departamento de Ingeniería Informática, Escuela Politécnica Superior, Universidad Autónoma de Madrid, C/Francisco Tomás y Valiente, 11, 28049 Madrid, Spain.
e-mail: pablo.varona@uam.es

Clinical olfactory tests are used to address hyposmia/anosmia levels in patients with different types of olfactory impairments. Typically, a given test is employed clinically and then replaced by a new one after a certain period of use which can range from days to several months. There is a need to assess control quality of these tests and also for a procedure to quantify their degradation over time. In this paper we propose a protocol to employ low-cost artificial noses for the quantitative characterization of olfactory tests used in clinical studies. In particular, we discuss a preliminary study on the Connecticut Chemosensorial Clinical Research Center Test kit which shows that some odorants, as sensed by an artificial nose, seem to degrade while others are potentiated as the test ages. We also discuss the need to establish a map of correspondence between human and machine olfaction when artificial noses are used to characterize or compare human smell performance in research and clinical studies.

Keywords: electronic noses, quality control of clinical olfactory tests, map of human and machine olfaction, olfaction, anosmia, olfactory dysfunction, artificial noses

INTRODUCTION

The assessment of olfactory function is an essential step for the diagnosis and treatment of olfactory disabilities. In particular, the quantification of the olfaction functional level is necessary to assess the recovery from or the progression of the smell dysfunction. The link between some types of smell impairments and the early detection of neurodegenerative (Albers et al., 2006; Barrios et al., 2007; Doty, 2008) or psychiatric (Atanasova et al., 2008) diseases also emphasizes the need of precise olfactory quantification procedures. Different protocols like the Connecticut Chemosensory Clinical Research Center (CCCRC) test (Cain, 1989), the University of Pennsylvania Smell Identification Test (UPSIT) (Doty et al., 1984), Sniffin' Sticks (Burghart, Wedel, Germany) (Hummel et al., 1997), T&T olfactometry, or the Odor Stick Identification Test for Japanese (OSIT-J) (Zusho, 1983) are employed throughout the world to assess the sense of smell in patients with different cultural backgrounds and disabilities. These tests consist of different odor containers that are presented to a patient to quantify specific odor thresholds, and to evaluate odor identification and discrimination.

A given test sample is typically employed in the clinic and then replaced by a new one after a certain period of use. In many cases the sample is opened, used, and closed again over an extended period of time which can range from days to several months. Thus, there is a need to assess control quality of these test samples and to establish a protocol to quantify their degradation over time.

Artificial noses are devices that use one or multiple sensors to mimic the sense of smell. The sensor reacts to an odorant and generates a signal that can be used for characterization, discrimination, or recognition purposes. The first artificial nose with a tin-oxide sensor was reported in the early 1980s. A wide range of sensor technologies have been developed since them to build many different types of artificial noses: electrical, gravimetric, optical, etc., [for a review of sensor technologies for artificial noses see (Stitzel et al., 2011)]. The applications of artificial noses cover areas such as safety, security, food, and beverage quality control, environmental monitoring, medical diagnostics, etc., (Dymerski et al., 2011; Stitzel et al., 2011).

In the following sections we propose the use of low-cost artificial noses to quantitatively characterize olfactory tests for clinical studies, and we discuss a preliminary study on the CCCRC kit. We also discuss the need to establish a map between human and machine olfaction for research and clinical studies.

MATERIALS AND METHODS

CONNECTICUT CHEMOSENSORIAL CLINICAL RESEARCH CENTER TEST

The CCCRC kit was developed by Cain and colleagues in 1988. This test consists of two parts: the butanol threshold test and the identification test. For the purposes of this study we will mainly concentrate on the butanol threshold test. The test was manufactured at Hospital Fundación de Alcorcón, Madrid with the collaboration of the Pharmacy Unit following the guidelines of the original article. The threshold test employs aqueous dilutions of 1-butanol where successive dilutions differ by a factor of three. The highest aqueous concentration is 4%. The number of dilution steps ranges from 0 to 8 depending on testing circumstances.

FIGURE 1 | The Connecticut Chemosensorial Clinical Research Center Test used in this study. The brown bottles contain the butanol threshold test, while the white plastic jars contain the identification test.

The test solutions were stored in 125 ml polyethylene bottles containing 60 ml of solution which are presented to patients during the test. The bottle closure has a pop-up spout that fits to both nostrils (see **Figure 1**). To sample a bottle, the person places the spout into both nostrils and then sniffs. There are two common ways to present the stimulus: the ascending method of limits (AMLs) procedure and the single staircase (SS) procedure. In the AML, the odorant and the water are presented sequentially from low to high concentration and the point of transition between no detection and detection is estimated. In the SS method, the concentration of the stimulus is increased following trials in which correct detection occurs. In both methods, the stimulus is presented in order from weak to strong. The SS procedure is more reliable and is used more often for threshold testing. Four correct choices in a row lead to end the test. The concentration at which this occurs marks the threshold.

Recently, Toledano and colleagues have described a short olfactory test based on CCCRC that predicts how severe is the smell loss (Toledano et al., 2009). This short test consists on determining the lowest concentration of butanol that the patient can detect four times in a row. Patients detecting beyond the butanol dilution number 3 have a high probability of normal olfaction.

The participant receives the *identification test* after the threshold test. The identification kit is composed of ten 180 ml opaque plastic jars containing 5 g of the substance in sachet-like packets of stimuli. Based on the performance of anosmic patients, we can ascertain that seven stimuli appeal almost exclusively to the sense of smell (baby powder, chocolate, cinnamon, coffee, mothballs, peanut butter, and bar soap) and one appeals to the trigeminal sense as well (Vicks). These eight items are presented in the same order to both nostrils. When presented with an item, the patient chooses from a 20-item list. The list contains the names of the eight test products and 13 distractors. In addition to the names on the list, responses of "no sensation" and "do not know" are permitted. The examiner provides corrective feedback whenever the participant makes an error. If the participant exhibits some evidence of function, but nevertheless makes mistakes, the examiner presents missed items for a second time. A correct answer upon a second presentation cancels a previous error. This allows the patient to rectify mistakes and thereby decreases the possibility of cognitive errors. In such cases, the first trial with an item serves as training. The score for the test comprises the number of olfactory items out of seven correctly identified and a notation regarding the ability to perceive trigeminal stimulation.

The outcome of the threshold and the identification tests is combined into a *composite score*, an average of the two tests. As previously defined (Toledano et al., 2003), a score of 5 points or more indicates normosmia, a score between 2 and 4.5 points indicates reduced olfactory function in terms of hyposmia, and a score of less than 2 points indicates functional anosmia.

To categorize progression of olfactory function, a subjective improvement and a normal score in butanol, threshold and composite, respectively, are regarded as a clinically significant improvement of olfactory function. Its correlation with previously established tests of olfactory function (e.g., the 12-item Cross-Cultural Smell Identification Test [CC-SIT, a subtest of the UPSIT] and the "Sniffin' Sticks" test) has been demonstrated in various studies (Kobal et al., 2000; Toledano et al., 2007).

ARTIFICIAL NOSES FOR QUALITY CONTROL OF OLFACTORY TESTS

During the last decades artificial nose technology has provided many successful examples of industrial applications [for recent reviews see (Dymerski et al., 2011; Stitzel et al., 2011)]. As this technology becomes cheaper and more accessible, the possibilities of potential use in medical and clinical studies have also increased. In this section we discuss a preliminary study on the use of a low-cost, portable artificial nose to quantitatively characterize the CCCRC test.

For the study reported in this paper we used an odor analysis platform capable of managing up to 16 samples with on/off control electrovalves and three auxiliary devices such as pumps, heaters, and/or mixers (see **Figure 2**). Although the platform can be equipped with a wide variety of sensor arrays, we aimed to study the performance of inexpensive sensors that could be used

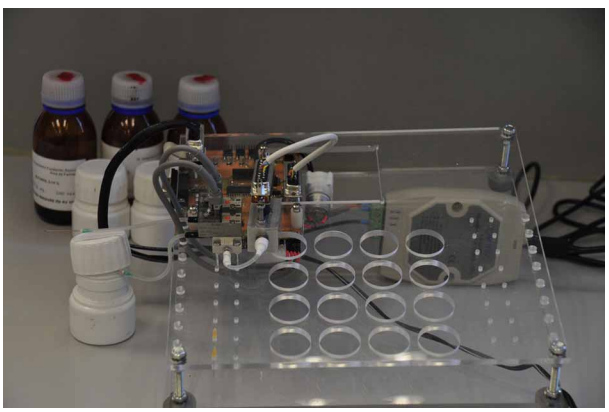

FIGURE 2 | Analysis platform used in this study with the TGS2600-B00 (Figaro Engineering Inc.) chemoresistive sensor.

in future widely distributed control quality devices. Thus, we selected the TGS2600-B00 (Figaro Engineering Inc.) chemoresistive sensor. This small sensor also provides a good combination of high sensitivity to low concentrations of odorants, long life, low power consumption, and robustness. The sensor is comprised of a metal oxide semiconductor layer formed on an alumina substrate of a sensing chip together with an integrated heater. In the presence of an odorant, the sensor's conductivity increases as a function of the concentration. A simple electrical circuit converts the change in conductivity to an output signal that corresponds to the odorant. The signal generated by the sensor is not amplified, it is only conditioned with an LMC6484 amplifier in voltage follower mode and sent to the Analog-to-Digital (A/D) converter module to be acquired and stored in a computer. The voltage range of this signal is 0–5 V.

From the control point of view, the main component used in this platform is an 8-bit PIC18F2455 microcontroller (Microchip Technology Inc.). This microcontroller is ideal for low-cost, low power (nano-Watt), and connectivity applications. The firmware was programmed with the MPLAB IDE development environment provided by the same company. The computer application software was developed with IDE Dev-C++ (GNU license) using Windows APIs. The equipment is connected to a PC via serial bus, by a proprietary protocol based on EIA-485. The transfer rate used in our study was 100 samples per second. For this application we used the rotary vane pump 01-K G-LC (Rietschle Thomas Company) in continuous vacuum mode (−20 mbar max.) running at 80% of its nominal voltage. All the instrumentation used in this study is characterized by its compactness and low price, which are important factors for a wide dissemination of quality control in olfactory tests.

Since the goal of our experiments is to assess the degradation of well-known odorants, we used the peak of the sensor signal to characterize the dilutions/odorants in a comparative manner (the peak corresponding to an aged olfactory test vs. the peak corresponding to a new test). The maximal separability between signals occurs typically at the peak. This is a simple and straightforward measurement that does not require further processing of the signal.

The performance of the TGS2600-B00 sensor has a moderate ambient temperature/humidity dependence as described in the manufacturer's worksheet (http://www.figarosensor.com/products/2600pdf.pdf). All experiments reported here were performed at a room temperature between 23°C and 24°C and a humidity of 30–35%. The acquisition of the signal was performed with the following protocol: the sensor was exposed to room air with no heating (OV) for 5 s, followed by 100 s at maximum heating (4.8 V) to clean the sensor (heater resistance is 83 Ω). After this cleaning phase, we took 5 s of absorption of the odorant (4.8 V for heating, this duration was chosen to avoid saturation), and finally 100 s of desorption (room air again at maximum heating). Five consecutive measurements of all odorants were taken following this protocol to avoid history dependent effects.

RESULTS

The quality control assessment of a clinical olfactory test can use the quantitative comparison of the artificial nose characterization of a brand new test and that of a test that has been used for a certain period of time. Here, we propose a simple quality analysis that consists on the quantitative assessment by the artificial nose of a dilution in the test as compared with the corresponding next dilution of a new test. A given olfactory kit fails the quality test when the response of the artificial nose for a dilution is close to the response obtained for the next dilution in the new test, which contains a lower concentration of the odorant.

As mentioned above, a short version of the CCCRC kit consists in the presentation of dilution number 3. Patients who cannot smell this dilution have a very high probability of olfactory dysfunction (Toledano et al., 2009). To illustrate and test the proposed quality protocol with the artificial nose, we will use dilutions number 2, 3, and 4 of the CCCRC threshold kit. Dilution 2 has a higher concentration of butanol than dilution 3, while dilution 4 has a lower concentration.

Figure 3 shows the response of the artificial nose to these dilutions corresponding to a one year old test (labeled as O for old) and to a new test (labeled as N). This figure illustrates that, after one year, the sensor signal that corresponds to the old dilution number 2 (2O) differs only slightly from the corresponding to the new dilution (2N) and remains far from the signal of the new dilution number 3 (3N). However, the signal corresponding to the old dilution number 3 (3O) is nearly as close to the signal corresponding to new dilution number 4 (4N) as to the new dilution number 3 (3N). Thus, as sensed by the artificial nose, dilution 3O could correspond to the next dilution (lower concentration) of a new test. In this situation, a quality control criteria based in the relative distance between the peaks of the signals would result in discarding the old test kit.

For each dilution, the maximum peak is reached a little after the stimulation is stopped. We remind that the protocol that we follow uses five consecutive measurements of the same odorant

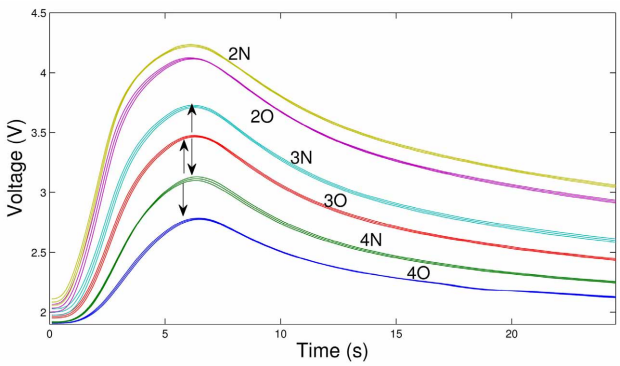

FIGURE 3 | Response of the artificial nose to three dilutions of the CCCRC kit for two different sets: a kit that has been used for over one year (dilutions labeled as 2O, 3O, and 4O) and a new one (dilutions labeled as 2N, 3N, and 4N). For each dilution we show the mean of five measurements (middle trace) together with the standard deviation (upper and lower traces). Note that the artificial nose signal for dilution 3O is nearly as close to 3N as to 4N. A quality control criteria based in the relative distance between the peaks of the signals would result in discarding the old test kit.

to avoid history dependent effects between different odorants. Note that we use the same desorption period for all of them and odorants that reach a higher maximum also start from a slightly higher baseline in this protocol.

We have also addressed the analysis with the artificial nose of the identification kit that the subject receives after the threshold kit in the CCCRC test. The purpose of this second kit is the evaluation of the ability to identify an odor. Thus, the analysis of the artificial nose in this case is oriented to determine the stability of the odorant over time. If the odorant changes significantly over the lifespan of the kit (typically one year), this could result in the lack of identification or in a wrong identification of the odorant by the subject. **Figure 4** shows the result of this analysis for two odorants (coffee and peanut butter). The aged tests (used clinically for more than one year) show a larger intensity as sensed by the artificial nose. The odorants in this identification kit are hidden under a gauze pad. The increase in signal amplitude observed in these aged jars is probably due to a degenerative chemical reaction or the presence of bacteria in the sample and in the pad because of the prolonged use of the kit. The rest of the odorants in the identification kit present a normal degradation as sensed by the electronic nose.

DISCUSSION

The precise quantification of olfactory perception is an essential step for the evaluation, diagnosis, and treatment of smell disorders and contributes to the success of a corresponding therapy. In the last decades several olfactory tests have been proposed to quantify odor thresholds and to evaluate human odor identification and discrimination. In this paper we have argued that olfactory tests could benefit from quality control performed by artificial noses to assess and guarantee their validity over their lifespan.

We have proposed a protocol to test the quality of a threshold clinical kit by comparing the response of the artificial nose to a given dilution of the test with the immediate next dilution of a new test. We have validated this protocol with the CCCRC threshold kit using an inexpensive chemoresistive sensor, which is a requirement for the successful dissemination of a quality control device for clinical purposes. The protocol does not need a complex processing of the sensor signal such as the use of classification or clustering algorithms required in many applications with artificial noses (Hierlemann and Gutierrez-Osuna, 2008; Muezzinoglu et al., 2010). As illustrated in **Figures 3** and **4**, the response of the sensor to the odorants in the CCCRC kit is highly reproducible. The comparative nature of the protocol (the results of the artificial nose on an old test are always compared with those of a new test) reduces the dependence of the results on sensor drift and environmental conditions such as temperature and humidity.

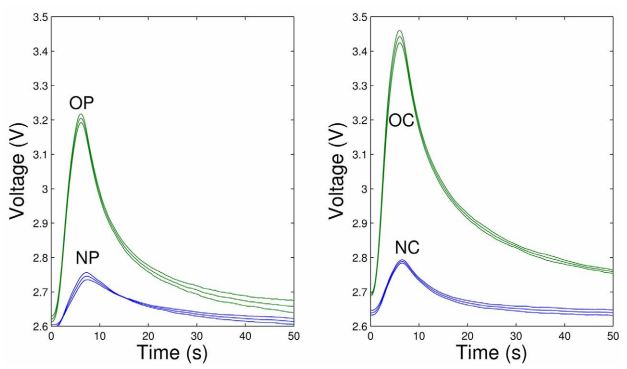

FIGURE 4 | Response of the artificial nose to two samples (peanut butter—left panel and coffee—right panel) of the identification part of the CCCRC kit. For each sample we show the mean of five measurements (middle trace) together with the standard deviation (upper and lower traces). Note that, as sensed by the artificial nose, the old peanut butter (OP) and coffee jars (OC) have higher odor intensity than the ones that contain a new test (NP and NC, respectively).

Human olfaction and machine olfaction have different mechanisms to sense smell, and thus the quantitative description of an odorant can also be very different for human and artificial noses (Burl et al., 2001; Schiffman et al., 2002; Lewis, 2004; Haddad et al., 2008, 2010). When artificial noses are used to characterize or compare human smell performance, the signals obtained with these devices need to be calibrated and matched to what is considered a normal perception by a human. In the context of clinical olfactory tests where the smell of odorants takes place under controlled conditions, the quantitative description of an odor by an artificial nose can be compared to an established reference to assess the quality of the test. The results shown in this paper indicate that this is possible, and that an inexpensive chemoresistive sensor can be used to characterize the threshold part of the Connecticut Chemosensorial Clinical Research Center Test. We have also reported that some odorants of the identification part of the CCCRC seem to be potentiated as sensed by the artificial nose while others tend to degrade as the test grows older. The analysis of these results suggests that the conditions to keep the odorants in the identification test could be improved.

Further efforts to build a map between human and machine olfaction can largely contribute to the assessment of odor perception and to develop novel protocols for odor impairment diagnosis and treatment. Artificial noses can be used to assess the intensity, steadiness, or temporal evolution of an odorant acting as a stimulus in a perception experiment. These experiments can be carried out with a standard feedback from the self-assessment of smell perception by the subject or through a more quantitative protocol that takes into account levels of brain activity in EEG (Lorig, 2000) and/or fMRI setups (Lombion et al., 2009) (which can largely contribute to a quantitative map between human and machine olfaction). In all cases, the assessment of odor temporal structure and intensity by artificial noses will improve the quantification and the standardization criteria in these protocols. In addition, artificial noses can also be used in closed-loop experiments with olfactometers and other stimulation devices to implement neural activity-dependent stimulation protocols.

ACKNOWLEDGMENTS

This work was supported by MICINN BFU2009-08473, TIN-2010-19607, and IPT-2011-0727-020000.

REFERENCES

Albers, M. W., Tabert, M. H., and Devanand, D. P. (2006). Olfactory dysfunction as a predictor of neurodegenerative disease. *Curr. Neurol. Neurosci. Rep.* 6, 379–386.

Atanasova, B., Graux, J., Hage, W. E., Hommet, C., Camus, V., and Belzung, C. (2008). Olfaction: a potential cognitive marker of psychiatric disorders. *Neurosci. Biobehav. Rev.* 32, 1315–1325.

Barrios, F. A., Gonzalez, L., Favila, R., Alonso, M. E., Salgado, P. M., Diaz, R., and Fernandez-Ruiz, J. (2007). Olfaction and neurodegeneration in HD. *Neuroreport* 18, 73–76.

Burl, M. C., Doleman, B. J., Schaffer, A., and Lewis, N. S. (2001). Assessing the ability to predict human percepts of odor quality from the detector responses of a conducting polymer composite-based electronic nose. *Sens. Actuators B Chem.* 72, 149–159.

Cain, W. S. (1989). Testing olfaction in a clinical setting. *Ear Nose Throat J.* 68, 316, 316–322, 328.

Doty, R. L. (2008). The olfactory vector hypothesis of neurodegenerative disease: is it viable? *Ann. Neurol.* 63, 7–15.

Doty, R. L., Shaman, P., and Dann, M. (1984). Development of the University of Pennsylvania Smell Identification Test: a standardized microencapsulated test of olfactory function. *Physiol. Behav.* 32, 489–502.

Dymerski, T. M., Chmiel, T. M., and Wardencki, W. (2011). The olfactory vector hypothesis of neurodegenerative disease: is it viable? *Rev. Sci. Instrum.* 82, 111101.

Haddad, R., Carmel, L., Sobel, N., and Harel, D. (2008). Predicting the receptive range of olfactory receptors. *PLoS Comput. Biol.* 4, e18. Public Library of Science. doi: 10.1371/journal.pcbi.0040018

Haddad, R., Medhanie, A., Roth, Y., Harel, D., and Sobel, N. (2010). Predicting odor pleasantness with an electronic nose. *PLoS Comput. Biol.* 6, e1000740. Public Library of Science. doi: 10.1371/journal.pcbi.1000740

Hierlemann, A., and Gutierrez-Osuna, R. (2008). Higher-order chemical sensing. *Chem. Rev.* 108, 563–613.

Hummel, T., Sekinger, B., Wolf, S. R., Pauli, E., and Kobal, G. (1997). "Sniffin" sticks': olfactory performance assessed by the combined testing of odor identification, odor discrimination and olfactory threshold. *Chem. Senses* 22, 39–52.

Kobal, G., Klimek, L., Wolfensberger, M., Gudziol, H., Temmel, A., Owen, C. M., Seeber, H., Pauli, E., and Hummel, T. (2000). Multicenter investigation of 1,036 subjects using a standardized method for the assessment of olfactory function combining tests of odor identification, odor discrimination, and olfactory thresholds. *Eur. Arch. Otorhinolaryngol.* 257, 205–211.

Lewis, N. S. (2004). Comparisons between mammalian and artificial olfaction based on arrays of carbon black-polymer composite vapor detectors. *Acc. Chem. Res.* 37, 663–672.

Lombion, S., Comte, A., Tatu, L., Brand, G., Moulin, T., and Millot, J. L. (2009). Patterns of cerebral activation during olfactory and trigeminal stimulations. *Hum. Brain Mapp.* 30, 821–828.

Lorig, T. S. (2000). The application of electroencephalographic techniques to the study of human olfaction: a review and tutorial. *Int. J. Psychophysiol.* 36, 91–104.

Muezzinoglu, M. K., Vergara, A., Huerta, R., and Rabinovich, M. I. (2010). A sensor conditioning principle for odor identification. *Sen. Actuators B Chem.* 146, 472–476.

Schiffman, S. S., Gutierrez-Osuna, R., and Nagle, H. T. (2002). "Measuring odor intensity with e-noses and other sensor types," in *Proceedings of the 9th International Symposium on Olfaction and Electronic Nose,* (Rome), 68–72.

Stitzel, S. E., Aernecke, M. J., and Walt, D. R. (2011). Artificial noses. *Annu. Rev. Biomed. Eng.* 13, 1–25.

Toledano, A., González, E., Onrubia, T. J., Herráiz, C., Mate, M. A., García, M., Navarro, M., Plaza, G., Aparicio, J. M., de los Santos, G., and Galindo, N. (2003). The Connecticut Chemosensorial Clinical Research Center olfaction test: values in healthy volunteers. *Acta. Otorrinolaringol. Esp.* 54, 678–685.

Toledano, A., González, E., Rodríguez, G., and Galindo, N. (2007). The validity of CCCRC test in patients with nasal polyposis. *Rhinology* 45, 54–58.

Toledano, A., Ruiz, C., Navas, C., Herráiz, C., González, E., Rodríguez, G., and Galindo, A. N. (2009). Development of a short olfactory test based on the Connecticut Test (CCCRC). *Rhinology* 47, 465–469.

Zusho, H. (1983). Olfactometry in Japan. *Rhinology* 21, 281–285.

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.

Received: 30 September 2011; paper pending published: 20 October 2011; accepted: 17 January 2012; published online: 02 February 2012.

Citation: Yáñez DJ, Toledano A, Serrano E, Martín de Rosales AM, Rodríguez FB and Varona P (2012) Characterization of a clinical olfactory test with an artificial nose. Front. Neuroeng. ***5****:1. doi: 10.3389/fneng.2012.00001*

frontiers in
NEUROENGINEERING

REVIEW ARTICLE
published: 04 January 2012
doi: 10.3389/fneng.2011.00019

Sensor selection and chemo-sensory optimization: toward an adaptable chemo-sensory system

***Alexander Vergara*[1]* and *Eduard Llobet*[2]**

[1] *BioCircuits Institute, University of California San Diego, La Jolla, CA, USA*
[2] *Department of Electronic Engineering, MINOS–EMaS, University Rovira i Virgili, Tarragona, Spain*

Edited by:
Ramon Huerta, University of California San Diego, USA

Reviewed by:
Eugenio Martinelli, University of Rome Tor Vergata, Italy
Stella Vallejos Vargas, University College London, UK

***Correspondence:**
Alexander Vergara, BioCircuits Institute, University of California San Diego, 9500 Gilman Dr., La Jolla, CA 92093-0402, USA.
e-mail: vergara@ucsd.edu

Over the past two decades, despite the tremendous research on chemical sensors and machine olfaction to develop micro-sensory systems that will accomplish the growing existent needs in personal health (implantable sensors), environment monitoring (widely distributed sensor networks), and security/threat detection (chemo/bio warfare agents), simple, low-cost molecular sensing platforms capable of long-term autonomous operation remain beyond the current state-of-the-art of chemical sensing. A fundamental issue within this context is that most of the chemical sensors depend on interactions between the targeted species and the surfaces functionalized with receptors that bind the target species selectively, and that these binding events are coupled with transduction processes that begin to change when they are exposed to the messy world of real samples. With the advent of fundamental breakthroughs at the intersection of materials science, micro- and nano-technology, and signal processing, hybrid chemo-sensory systems have incorporated tunable, optimizable operating parameters, through which changes in the response characteristics can be modeled and compensated as the environmental conditions or application needs change. The objective of this article, in this context, is to bring together the key advances at the device, data processing, and system levels that enable chemo-sensory systems to "adapt" in response to their environments. Accordingly, in this review we will feature the research effort made by selected experts on chemical sensing and information theory, whose work has been devoted to develop strategies that provide tunability and adaptability to single sensor devices or sensory array systems. Particularly, we consider sensor-array selection, modulation of internal sensing parameters, and active sensing. The article ends with some conclusions drawn from the results presented and a visionary look toward the future in terms of how the field may evolve.

Keywords: electronic nose, metal-oxide gas sensors, sensor optimization, sensor-array optimization, tunable sensors, active sensing

INTRODUCTION

The idea to mirror the biological senses, particularly the biological sense of olfaction, with artificial electronic systems has been a human dream for many years (Persaud and Dodd, 1982; Gardner and Bartlett, 1999). The choice of olfaction is not coincidental. While for humans, whose vision and hearing senses are their primary mode of communication with the outside world, olfaction is a rather little used sense, as demonstrated by the relatively normal lives led by people with anosmia – people who cannot smell – for most of the animals, olfaction is the primary means of exploration and communication (Bhandawat et al., 2007). The biggest challenge in performing such an imitation, though, is how to reliably emulate this system by understandable artificial mechanisms, commonly referred to as electronic noses (e-nose). e-Nose, or their odorant chemo-receptors to be more precise, play an important role in this challenge not only because they serve as oversimplified, yet accurate, reproductions, and simulations of the biological sense of olfaction, but also because these artifacts are able to non-invasively detect, acquire, interpret, select, and organize the sensory information of certain situations that humans can not perceive or understand (Persaud and Dodd, 1982; Freund and Lewis, 1995; Dickinson et al., 1996; Gardner and Bartlett, 1999). The capabilities of odor chemo-sensors are broad and include many challenging tasks such as discriminating organic compounds with chain lengths that differ by a single carbon atom (White et al., 1996; Persaud and Travers, 1997). However, their limitations in characterizing odor-stimuli, including the poor sensitivity to analytes and the lack of reproducibility in their responses in repeated trials, are still very serious (Moseley and Tofield, 1987). With the advent of the latest technological breakthroughs, recent progress in chemical micro-sensory systems has been stimulated by inter-disciplinary perspectives at the intersection of materials science, micro- and nano-technology, microelectronics, and signal processing/pattern recognition in an attempt to ameliorate these apparent limitations. For example, the effect of microstructure, size feature of the material, the advantages of nano-structured materials (e.g., nanowires, nanorods, nanoribbons), and the efforts to functionalize sensing materials by adding catalysts have been widely studied from the

material perspective to improve the sensors' sensitivity and selectivity as well as to reduce their power consumption and time of response (Yamazoe, 2005; Franke et al., 2006; Comini et al., 2009; Gurlo, 2010; Stoycheva et al., 2011). On the other hand, different feature selection, feature extraction, and pattern-recognition techniques, from the signal processing and machine learning viewpoint, have also been implemented with remarkable results in many critical applications (Di Natale et al., 1995; Eklöv et al., 1997; Nakamoto et al., 1997; Gutierrez-Osuna et al., 1999; Muezzinoglu et al., 2009; Vergara et al., 2011). However, while maintaining and expanding this fruitful interaction is essential for solving these problems ahead, this collaborative platform will be lacking a key player until the material and hardware development component improves, since the current sensory modalities available do not meet the power consumption and dimension restrictions required for particular real-time applications. Therefore, among the many different strategies implemented in the literature, interacting with the conditioning parameters at the sensor level (e.g., working temperature for metal-oxide gas sensors) is the only viable solution to overcome these annotated problems (Moseley and Tofield, 1987).

In principle, almost every odorant-sensing technology offers the advantage of being tunable through the selection of parameter values. Interacting with such parameters influences many critical qualities of the measurement, including sensitivity to analytes and reproducibility. For example, the temperature-modulation technique for metal-oxide gas sensors, takes advantage of such a relationship to enrich the information content of the sensor, since it directly alters the reaction kinetics at the sensor surface in the presence of an odorant (Sears et al., 1989a,b, 1990; Nakata et al., 1992, 1996, 1998; Semancik and Cavicchi, 1999). A thorough understanding of how such interactions take place in the chemo-sensory system requires quantitative characterizations of the response of individual sensors, both within and among chemical stimuli. This approach will enable us to generalize the relationship between the control variable and the target quality.

Once the odorant-sensing/parameter interaction is known (or can be inferred from previous observations), a natural venue to follow is the optimization of the chemo-sensory system with respect to a solid criterion that properly expresses the observed goal. A number of approaches under the notion of optimization have been implemented in the literature, but only a handful of authors have approached the problem in a systematic fashion. In this context, the purpose of this paper is to provide the reader with a critical review of the different efforts that have been made in the context of sensor selection and sensor optimization in the chemo-sensing community for the last years. We will be visiting these criteria individually as we proceed further in this review.

Before embarking upon the subject of optimization, we would like to make a final point in the context of terminology. The label "optimization" has been very popular for many years to describe the terms of "sensor optimization" and "sensor-array optimization" interchangeably. We believe, however, that this terminology can be very misleading, since the former mostly refers to finding the "optimal operational condition" of the sensor device, whereas the latter usually means selecting an "optimal" combination of sensors between a potentially large pool of different sensors that are best suited to the identification task – pretty much as feature selection. Therefore, although these two groups of procedures are fully complementary and valuable tools to analyze our chemical sensors, they deserve to be treated as separate topics. For this reason, we organize the structure of this review according to which of these two aspects of "adaptation" (or sensor optimization among many other names considered) is taking place at the sensor level. It is not say that this adaptation can occur only at one or another side of the coin; sensor optimization may also involve coordination of other aspects, such as the adaptability of the chemical sensors at multiple levels and time scales through their operating parameters targeted when the environmental conditions change. Accordingly, we have decided to give three more specific threads to run this review. The first is gain control in those methods that have been implemented to optimize the chemical sensors when used as a sensor array. The second, which is the other side of the same coin, is a functional analysis of the coding schemes used in the optimization of each chemical sensor individually. And the third is a systematic analysis of the advantages derived from the coding scheme used by the early optimization stages implemented on chemical sensors in an active fashion, or the so-called adaptive sensing optimization. Finally, in order to gain a better understanding of how the optimization processes are occurring at the sensor level, or even to discuss whether the processes are or are not working, we have included an overview of the operating temperature dependence for the response of semiconductor gas sensors at the beginning of the Section "Metal-oxide Gas Sensors and Their Operation: Initial Optimization Methods." We strongly recommend people who may not be familiarized with the functioning process of chemical sensors, specifically metal-oxide gas sensors, to review this section.

METAL-OXIDE GAS SENSORS AND THEIR OPERATION: INITIAL OPTIMIZATION METHODS

METAL-OXIDE GAS SENSORS AND TEMPERATURE DEPENDENCE

Undoubtedly, metal-oxide gas sensors have become one of the most widely used sensing technologies in machine olfaction for a diversity of applications. Because of the strict and highly deterministic dependence of the sensor response on its operating temperature, governing the basic operating principle of this sensing technology, we believe that it is worthwhile to provide a good insight into the dynamic behavior and operating principle of metal oxide sensors so we can get a better understanding of how the optimization processes described here take place at the sensor level. Accordingly, the followings lines of this section feature the basic operating principle of the said chemo-sensing technology. The sensitive layer, in this case the metal-oxide film that is made of particles that range from nanometers up to microns, possesses two operating mechanisms. The former is associated with an ideally specific interaction of the surface with the target analyte, whilst the latter refers to an effective transduction of the bulk conductance. If the interaction takes place exclusively at the surface of the sensitive layer, then the bulk conductivity does not contribute and represents only a shunt which decreases the signal-to-noise ratio. On the other hand, for materials in which interaction originates in the bulk of the sensitive layer, the response of the sensor, and in particular its time of response, is affected by the thickness and porosity of the material, giving faster response times for thin films than for thick films (Yamazoe, 2005; Franke et al., 2006; Stoycheva

et al., 2011). Accordingly, for a given type of base material, the sensor property sensitively depends on the structural features, the presence and state of catalytically active surface dopants, and the working temperature.

The central reaction mechanism responsible for most of the chemical compound responses/interactions involves changes in the concentration of surface oxygen species, such as O^{2-} (Göpel, 1985, 1988; Göpel et al., 1991). The formation of such ions means that the oxygen adsorbed at the gas/solid interface abstracts electrons from the conduction band of the sensing material, which results in the development of Schottky potential barriers at the grain boundaries. In the case of an *n*-type semiconductor such as SnO_2, on the one hand, the electrons come from ionized donors via the conduction band, the charge-carrier density at the interface is reduced, and a potential barrier to charge transport, $\Delta\phi$, develops. At the junctions between the grains of the solid (see **Figure 1**), the depletion layer and the associated potential barrier are responsible for high resistance contacts which dominate the resistance of the sensor. Thus, depending on the temperature, oxygen is ionosorbed on the surface predominantly as O_{2-} ions below 150˚C, or as O^- ions between 150 and 400˚C, which is the general operating temperature range. Above 400˚C, the parallel formation of O^{2-} occurs, which is then directly incorporated into the lattice above 600˚C (Barsan et al., 1999). In the case of a *p*-type oxide, on the other hand, adsorbed oxygen acts as a surface acceptor state, abstracting electrons from the valence band, and hence giving rise to an increase in the charge-carrier (holes) concentration.

In response to an analyte and under stationary conditions (i.e., without humidity, constant flow, and fixed operating temperature), the sensor involves an exponential change in the conductance/resistance across its sensing layer. This resulting change can be interpreted as a shift of the state of equilibrium of the surface oxygen reaction due to the presence of the target analyte, which can be either a reducing or oxidizing specie. The response of semiconductor gas sensors to reducing species implies a change in the concentration of adsorbed oxygen species. On the other hand, oxidizing species can interact with the sensor surface in a variety of ways; for example, interacting directly with the surface and forming negatively charged ionosorbed species or in competition with ionosorbed oxygen or oxygen ions for the adsorption sites available (Ruhland et al., 1998). These changes modulate the height of the potential barriers and thus the conductance of the sensing layer. The reason these characteristic conductance–temperature profiles arise is summarized as follows:

- There are different adsorbed oxygen species such as O^-, O^{2-}, and O_{2-} over the temperature range (Sears et al., 1989a; Barsan et al., 1999).
- Different gases have different optimum oxidation temperatures (Moseley and Tofield, 1987).
- Adsorption, desorption, and diffusion rates (of oxygen species, reducing and oxidizing gases, and oxidation by-products) are temperature-dependent (Clifford and Tuma, 1983; Nakata et al., 1991; Wlodek et al., 1991).

Accordingly, when the operating temperature of the sensor varies, the kinetics of adsorption, desorption, and reaction occurring at the sensor surface in the presence of atmospheric oxygen and other reducing or oxidizing species is altered. This approach leads to sensor responses (e.g., transient conductance patterns) that are characteristic of the species present in the gas mixture. Having such an easy way of interacting with the sensor and its characteristics justifies the use of temperature as a control variable in a deterministic setting and, thus, the optimization of the sensor device with respect to this parameter (i.e., the sensor's operating temperature).

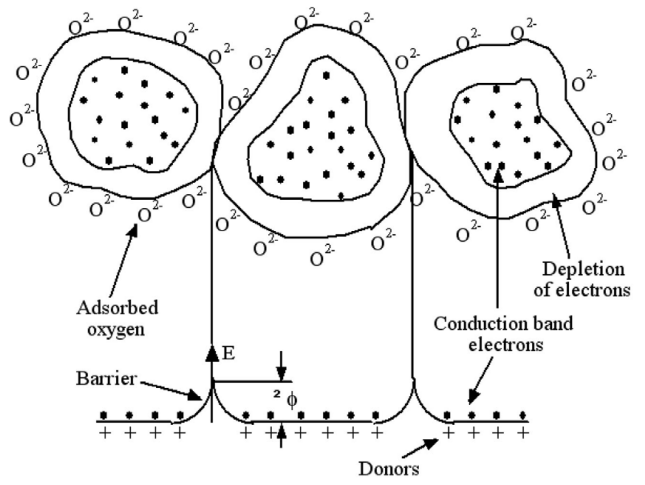

FIGURE 1 | Structural and band model showing the role of intergranular contact regions in determining the conductance over a polycrystalline metal-oxide semiconductor. Three grains with adsorbed oxygen providing surface depletion layers. The depleted layers are responsible for a high contact resistance. For conduction, electrons must cross over the surface barriers.

INITIAL METHODS

More than two decades have passed since Parsaud's pioneering publication on machine olfaction and e-nose appeared in literature (1982). Researchers have since devoted their work to developing signal processing procedures and optimization strategies to ameliorate the performance of metal-oxide gas sensors. This section presents an introduction to such first deployed optimization methods. One of the pioneering works was presented by Corcoran et al. (1998), who applied a triangular waveform (4.16 mHz) to modulate the operating temperature of eight commercially available gas sensors (from Figaro Engineering Inc., Japan, http://www.figaro.co.jp) between 250 and 500˚C. They extracted features from the sensor transients using two different approaches. The first approach consisted of sub-sampling the response transients obtaining a 26-point vector (equivalent to 10˚C steps) per transient. The second one, on the other hand, consisted of calculating eight secondary features from each response transient, such as the time to maximum value, time to minimum value, maximum positive slope, etc. They implemented then an optimization procedure to determine which sensors and features should be used to better classify the aromas from three loose leaf teas using a neural network classifier. The optimization process

consisted of applying genetic algorithms (GAs) for variable selection (Davis, 1991). By applying this technique, the authors showed that it was possible to reach a high success rate in tea classification (93%) using only 21 dynamic features out of the 208 available features. However, the optimization of the temperature-modulating signal (frequency, temperature range, and waveform type) was not considered at the time by the authors. Had this optimization been envisaged, further improvements in classification performance would have been obtained.

More recently, Fort et al. (2002) and Fort et al. (2003) used temperature-modulated metal-oxide gas sensors (sensors modulated with a pure sinusoidal signal) to show that the selection of the signal frequency was of paramount importance for gas identification. The authors demonstrated that if the temperature of the sensors is varied relatively quickly in comparison with the chemical response time, the sensor resistance varies as a function of the temperature with an exponential law characteristic to metal oxides. AS a result, the response shape would have only a slight dependence on the chemical environment. On the other hand, when the operating temperature varies slowly enough in comparison to the chemical response time, the response profile gives a series of quasi-stationary chemical responses. The best discrimination among the species studied (vapors from water solutions containing ethanol and other volatile organic compounds) can be obtained then by selecting a temperature profile with a period close to that for the chemical response time of the sensor. Accordingly, these results suggested that the effectiveness of the temperature-modulation analysis depends on the period of the sine wave, which must be chosen in agreement to the chemical reaction rate of each sensor. Around the same time, Choi et al. (2002) and Huang et al. (2003) got similar results that the ones presented above. The authors evaluated the effect of utilizing different kind of temperature modulation signals, such as pulse, trapezoid, triangular, and saw-tooth, as well as modulating frequency values in the sensor response performance. In particular, utilizing different modulating frequencies values (50, 30, 40, and 20 mHz), Huang et al. (2003) experimentally demonstrated that different and more odor specific response patterns were developing in the sensor response as the modulating frequency was taken nearer to low values (e.g., 20 mHz), suggesting that low-frequency temperature modulation signals are more relevant to alter the kinetics of adsorption, diffusion, and reaction phenomena (i.e., the interaction of odorous compounds and the gas-sensitive surface) occurring at the sensor surface.

SENSOR-ARRAY OPTIMIZATION

A wide variety of sensors, feature extraction, and feature selection methods that are available to the experimenter when considering a new sensing problem have been described elsewhere in the literature (Guyon and Elisseeff, 2003; Rodriguez-Lujan et al., 2010; Vergara and Llobet, 2011). However, when one is working with an array of non-specific sensors, the biggest concern to the experimenter is the number of sensors needed to form a chemo-sensory array. One manner to approach this dilemma might be to augment an existing array by adding sensors appropriate to the new task. However, this is actually a computationally expensive and potentially wasteful solution because using more sensors does not necessarily guarantee improvement in the overall performance. As a consequence, the most conceivable way to address this issue is to design an optimal array of sensors (even comprising completely different sensing technologies) that would promote the maximum accuracy with which the sensory system can estimate the stimulus or optimally discriminate between neighboring stimuli. A number of theoretical studies have been performed with the notion of "array optimization" and are now available in the literature, where can be explored when the experimenter is approaching to a new odor identification scenario. One of the pioneering investigations in this context is the one presented by Zaromb and Stetter (1984), who proposed, over 20 years ago, a theoretical model to estimate the minimum number of parameter $P = (S$ sensors $\times M$ operating modes) that would be required to discriminate a mixture of up to A analytes from a pool of n different odorants. By assuming that the response of each sensor is noiseless and binary related to each odor stimulus (i.e., response/no response), they argued for a combinatorial measure of the number of sensors required to detect a given number of chemical species as

$$2^p - 1 \geqslant \sum_{i=1}^{A} \frac{n!}{(n-i)!i!} \tag{1}$$

The authors, whose work was subsequently corroborated by Alkasab et al. (2002), proposed a "rule of thumb" in their work according to which sensors and operating modes should be selected so that each of the P parameters does not respond to more than P/A individual compounds. Later, in a seminal paper presented by Niebling and Müller (1995), the authors proposed the use of an inverse feature space to design sensor arrays. In this inverse feature space, each of the n analytes was represented as a separate dimension, and each of the s sensors was then represented as a point in this n-dimensional space. They showed that this visual representation should enable the experimenter to detect potential discrimination problems and to design new sensors to adequately address these problems. Gardner and Bartlett (1996) proposed a computational model for cross-selective sensors that also considers the effects of noise and errors. By using the ratio between the total volume of the sensor space and the volume made up by the sensor errors, the authors estimated an upper limit of the number of analytes that can be discriminated by the given array. As a final result, they proposed a measure of performance, which was essentially equivalent to the classical Fisher's linear discriminant analysis (LDA) ratio (i.e., the ratio of between-class distance to within-class variance).

It was not until the early 2000s when Pearce and Sanchez-Montañes (2003) implemented for the first time an information-theoretic approach for the optimization of chemo-sensory array systems. In particular, they demonstrated how the "tunings" of individual sensors may affect the overall performance of the array. In order to demonstrate the effects of noise and tuning on array performance, they incorporated the concept of "hyper-volume of accessible sensor space" (V_S), a volume in the sensor space that contains the sensor-array response to a specific set of analytes.

As **Figure 2A** suggests, for a three-odor by two-sensor problem, the collinearity limits the number of possible sensor responses. Therefore, the maximum number of analyte mixtures that can be discriminated by the array is limited by the ratio between V_S and V_N (the hyper-volume defined by the accuracy of the sensor array; see **Figure 2B**).

Assuming that the errors/noise ratio does not exhibit any correlation with the analyte stimulus, the authors showed then that the geometric interpretation in **Figure 2** can be expressed by means of the Fisher information (FI) matrix defined as

$$J_{i,j}(c) = \int p(x|c) \left(\frac{\partial}{\partial c_i} \ln p(x|c) \right) \left(\frac{\partial}{\partial c_j} \ln p(x|c) \right) dx, \qquad (2)$$

where c is a vector containing the concentration of the analytes, x is the response of the sensor array to the stimulus c, and $p(x|c)$

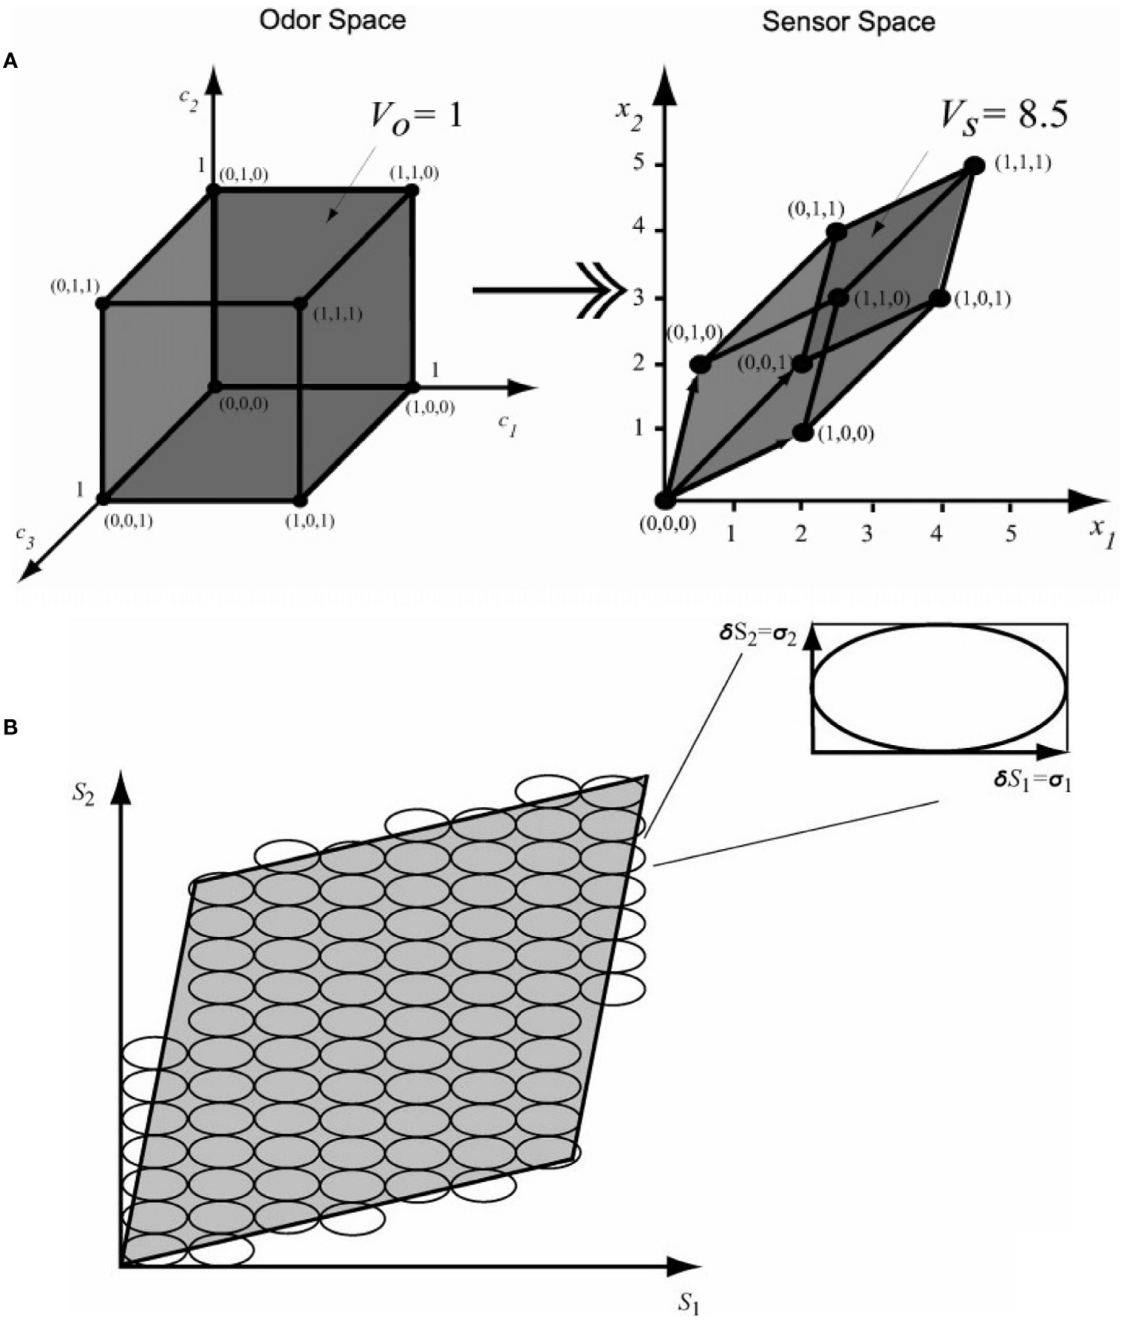

FIGURE 2 | (A) Visualization of a three-odor-to-two-sensor transformation. **(B)** The maximum number of feature vectors that can be discriminated is the ratio between the hyper-volume of the accessible sensor space (V_S) and the accuracy of the sensor-array response. Figure reproduced with permission from Pearce and Sanchez-Montañes (2003) Copyright 2003 Wiley-VCH, Weinheim.

is the conditional probability of observing the sensor response x upon a given stimulus c. FI is important in this context because it provides a lower bound (i.e., best-case case) on the accuracy with which the stimulus, c, can be predicted from the sensor response x. This lower limit is determined as

$$\mathrm{var}\left(\hat{c}|c\right) = \sum_{i=1}^{S} \left(J^{-1}\left(c\right)\right)_{ii} \tag{3}$$

where "*var*" means the variance, and $\hat{c}$ is the estimation/prediction of the component i of c, $i = 1, \ldots, N$. This result was called the "Cramér–Rao Bound" and limits the performance of the best unbiased estimator that can be built.

In order to use these theoretical constructs in practice, the first stage one should perform is the formal description of the sensory context C and a clear specification of the task. The context C, on one hand, quantitatively describes the likelihood of occurrence of each odor stimulus, whereas the chemo-sensory task, on the other hand, is an interpretation of the sensory response (i.e., a quantification or identification task). Once the sensory context and task are properly defined, one then would assume a parametric density $p(x|c)$ for each individual sensor, estimate the parameters from experimental data (i.e., by measuring the sensor-array responses to a number of analyte mixtures), compute the FI using Eq. 2, and finally compute the expected accuracy of the array using Eq. 3. This resulted accuracy estimate would then be used as a "figure of merit" to select an optimal array configuration from a pool of cross-selective sensors. Once the "optimal" array condition is found, a catalog of parameters for each sensor used within practical systems today may then be envisaged, which would make the optimization of sensory array systems to particular detection tasks a simple routine operation.

More recently, in the same domain of information theory, Muezzinoglu et al. (2010) introduced a sensor-array optimization scheme for odor identification. The authors demonstrated the effects of tuning the sensor's operating parameter in a chemo-sensory array by incorporating a measure-index widely used in signal theory, namely the Mahalanobis distance (MD), which gives a quantitative measure of the separability among probability distributions – odor classes for this specific machine olfaction application case. Since the chemo-sensory records associated to a given odor class have a certain variability regardless of the features selected, they can be assumed to be probability distributions that log the history of each sensor in response to a specific odor class over a feature space. Therefore, optimizing this index over a controllable operating parameter (e.g., the operating temperature in metal-oxide gas sensors) of the sensory device, would result then in improving the classificatory capabilities of the sensor itself, i.e., maximizing the spread of the class prototypes (the class centers) in the feature space while the response variability within each class is minimized.

To demonstrate their scheme, the authors first assumed a two-odor class formulation, where all the possible measurements may belong to one of the two disjoint classes C_1 and C_2 in a specific feature space. Then, given a sample x^s, their goal was to accurately determine which of the two-class-conditional distributions $f(x|C_1), f(x|C_2)$ was more likely to have produced x^s. The squared MD between two-class-conditional distributions is given by

$$D^2\left(C_1, C_2\right) = \left(\mu_1 - \mu_2\right)^T S_{1,2}^{-1}\left(\mu_1 - \mu_2\right), \tag{4}$$

where $\mu_i = \langle x|C_i\rangle$, $i = 1, 2, \ldots$, are the class centers and $S_{1,2}$ is the weighted average of the two covariance matrices S_1 and S_2 associated to the two-class-conditional distributions. For normally distributed classes, the MD index, which is proportional to the distance between-class centers (the between-class scatter) and inversely proportional to the individual co-variances (the within-class scatters), constitutes the best-possible quantification of the overlap. In this particular odor-discrimination case, this index also becomes the most accurate indicator of the classification performance for any unbiased classifier in the sense that the probability of misclassification is in inverse proportion with the MD value.

In a more generic case, i.e., when the number of classes is greater than two, the between-class scatter component of MD that promotes the dispersion of class centers can be generalized by the sum of their pair-wise distances, thus,

$$\mathrm{MD}^2 = \sum_{i,j=1}^{|C|} D^2\left(C_i, C_j\right) \tag{5}$$

where $|C|$ denotes the number of classes in the problem.

Since the class-conditional distributions are originally unknown to the designer, it is important to be able to estimate the MD index value from previous observations, i.e., previous measurements. Being dependent on the mean and variance, the sample MD is obtained by substituting these two moments by their sample estimates:

$$\mathrm{MD}^2 = \sum_{i,j=1}^{|C|} \left(\hat{\mu}_i - \hat{\mu}_j\right)^T \hat{S}_{i,j}^{-1}\left(\hat{\mu}_i - \hat{\mu}_j\right) \tag{6}$$

where $C = \{1, \ldots, |C|\}$ denotes the class labels, from which each $i \in C$ class is represented by n_i pre-recorded samples. Each class center $\hat{\mu}_i$ is approximated by the sample average of all samples in class i. The joint covariance $S_{i,j}$ is given by

$$\hat{S}_{i,j} = \frac{n_i \hat{S}_i + n_j \hat{S}_j}{n_i + n_j - 2} \tag{7}$$

being $\hat{S}_i$ the sample covariance matrix, i.e., the average of the outer products of the observations in class i.

Intuitively, the resulting MD index quantifies the difficulty of the classification problem. When this quantity is large, an arbitrary classifier is expected to perform with higher accuracy, since, relatively to a small MD, the distribution within each class is shrunk (the within-class scatter becomes small) and the two classes are located away from each other in the feature space (i.e., the between-class scatter is large).

Being θ a parameter of a sensor array that alters the sensor response characteristics, the problem configuration is then

expected to be sensitive to the said operating parameter, making thereby the MD index dependent on θ. Hence, the value

$$\theta^* = \arg\max_{\theta} \mathrm{MD}^2(\theta) \quad (8)$$

defines an optimum operating condition for the classification problem at hand.

Although this optimization criterion is applicable to any number and complexity of probability distributions, hence, to any type and number of odorants as well as any sensor technology with a conditioning parameter, the authors have shown the applicability of their approach to a particular three-class classification problem, i.e., ethanol, acetaldehyde, and ammonia. **Figure 3A** shows the maps of 30 samples, grouped with respect to their class labels, to the selected feature space for heater voltage values applied to the sensor array. As this example illustrates, the three classes move and change their relative positions with the temperature, making the classes easier or more difficult to separate along the sweep. Note also that the two-sensor responses can be highly correlated at certain temperatures and uncorrelated at others. As a measure of separability, the evaluation of the MD estimate, given in Eq. 8 for the triple of classes at each operating temperature labeled as the parameter θ, yields the profile shown in **Figure 3B**. Based on this evaluation, the best operating condition to distinguish among the triple of classes is determined by the maximum of this curve, which occurs at θ = 5.4 V, i.e., the best voltage applied to the sensor heater that yields the optimal sensor's operating temperature.

Utilizing a similar statistical argument, Raman et al. (2009) pioneered the development of an optimization method to design micro-sensing arrays for complex chemical sensing tasks. The method consisted of utilizing statistical methods to systematically assess the analytical information obtained from the conductometric responses of chemo-resistive elements at different operating temperatures, i.e., the similarity/orthogonality of responses; test their reproducibility; and determine an optimal set of material compositions to be incorporated within an array of sensors for the recognition of individual species. They presented qualitative and quantitative approaches to determine both the sufficiency of the chosen materials for sensing targets in the test matrix and an optimal array configuration for the desired application.

In order to optimize the array configuration, the authors presented a modular approach, a sophisticated temperature pattern (see **Figure 4A**) and micro-hotplate platforms with different metal-oxide chemo-resistors (see **Figure 4B**) that examines the target matrix with five high-priority chemical hazards (i.e., ammonia, hydrogen cyanide, chlorine, ethylene oxide, and cyanogen chloride). The temperature program used to operate the sensing elements toggles the temperature between (a) 32 ramp values that sample most of the temperature range of the device and (b) four different baseline temperature values to allow relaxation

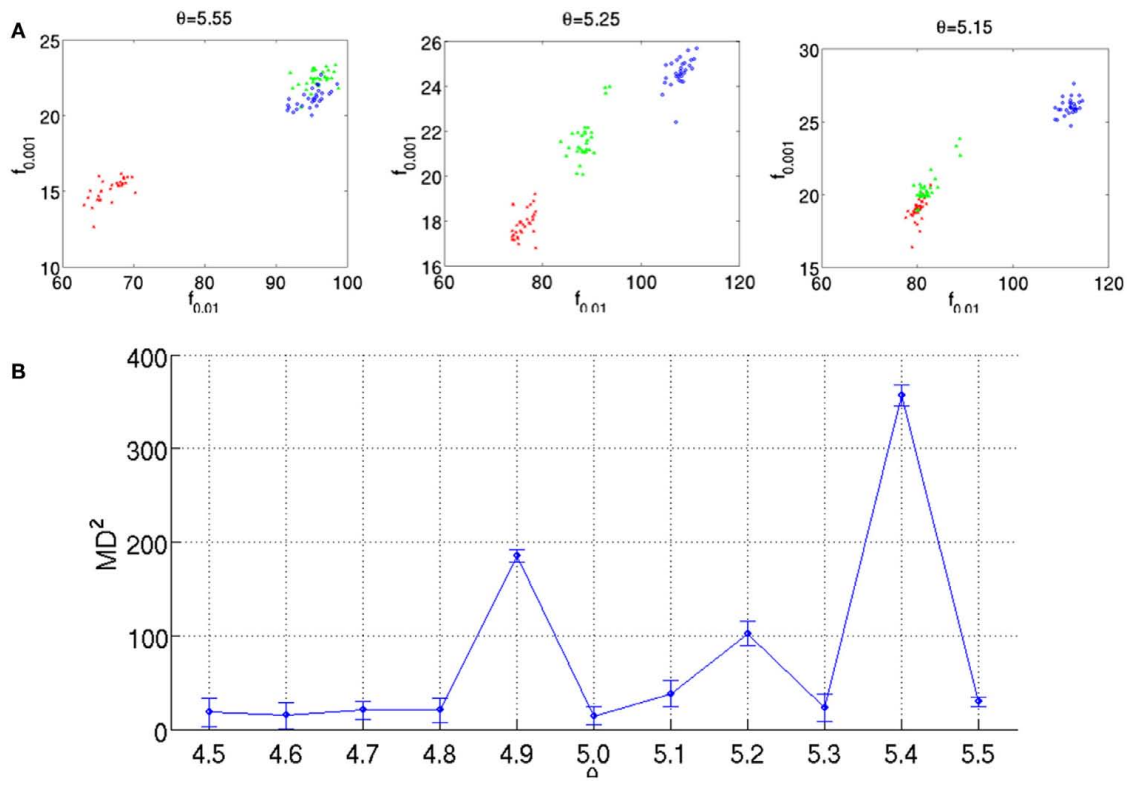

FIGURE 3 | (A) Feature maps obtained for three operating temperatures as indicated by θ on each figure. The two features (i.e., axes) used in this representation are the transient features extracted from the response x_1 of TGS2600 and x_2 of TGS2610. Each of the analyte classes contain 10 samples, which are labeled with a different color/shape on the maps. **(B)** MD2(θ) evaluations estimated from the dataset for 11 heater voltages. Figure reprinted from Muezzinoglu et al. (2010), Copyright 2010, with permission from Elsevier Science.

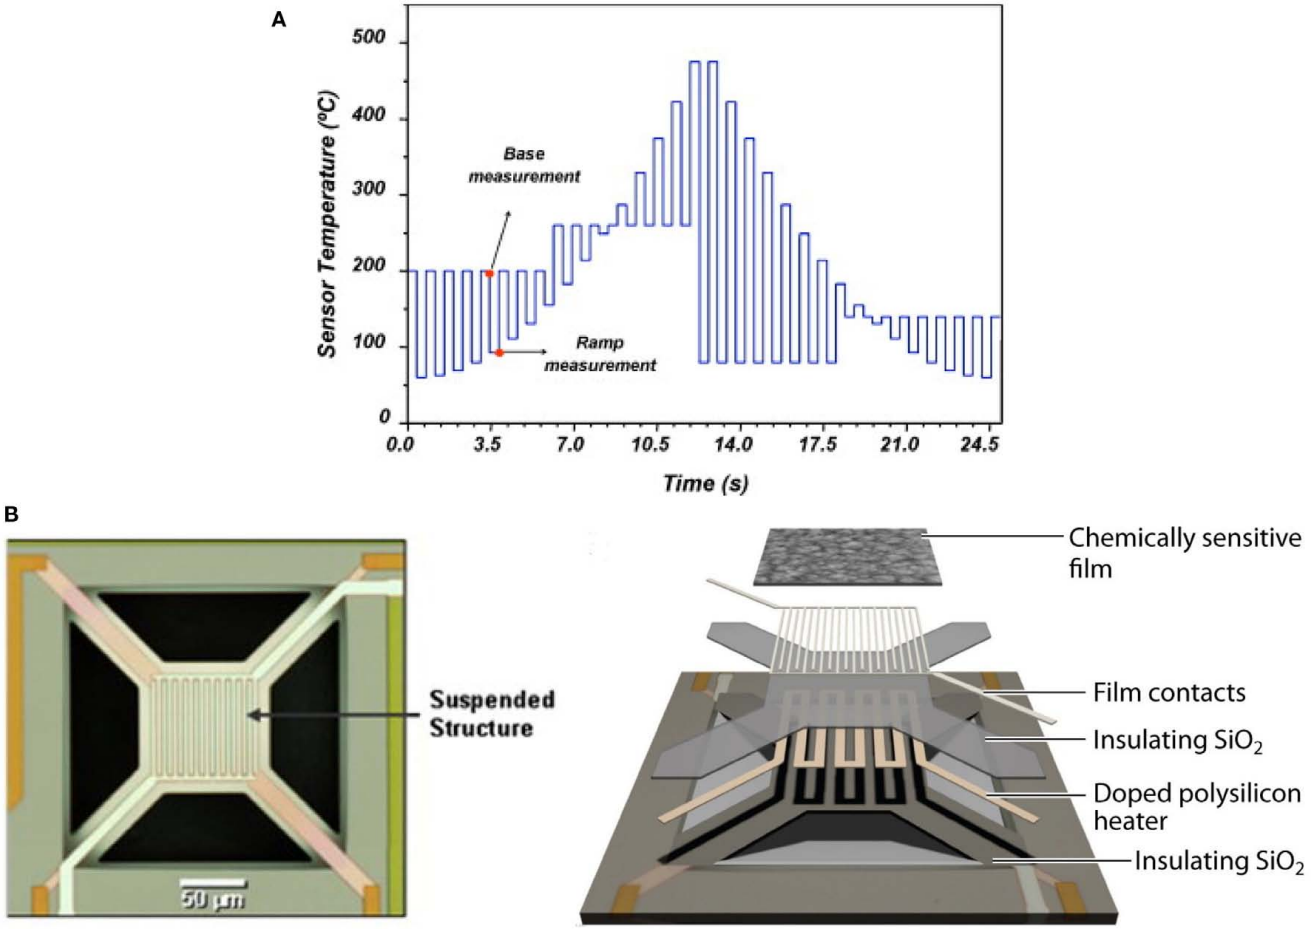

FIGURE 4 | (A) Temperature program used for the detection of high-priority chemical hazards. A conductance measurement was made at each base and ramp temperature, but only the ramp values were used for further analysis in this study. **(B)** Micro-sensor array platform. A layered schematic showing the three primary components of the micro-sensor elements: polycrystalline silicon heater, interdigitated platinum electrodes, and metal-oxide sensing film. Figure reprinted from Raman et al. (2009), Copyright 2009, with permission from Elsevier Science.

toward some initial state prior to each ramp temperature. These different baselines allow thus different film–analyte interactions (i.e., adsorption/desorption, decomposition, and reaction) at the sensing surface prior to the ramp measurements. Then, they defined the following objective function with three components:

$$O = \gamma_1 J - \gamma_2 N_1 - \gamma_3 N_2 \tag{9}$$

where J is the maximization term that takes into account the sufficiency of solution (i.e., separability of the five chemical analyte clusters from different background conditions and from each other); N_1 and N_2 are two penalty terms that symbolize the number of different materials used and the array size respectively; and γ_1, γ_2, and γ_3, are component weights. The two penalty terms allow comparison between solutions with different numbers of materials and array sizes. In order to be able to increase the objective function, each new material or array element must increase the analyte's cluster separability sufficiently to compensate for its "cost." This cluster separability is derived from Fisher's LDA as follows:

$$J = \frac{\text{trace}(S_B)}{\text{trace}(S_B) + \text{trace}(S_W)} \tag{10}$$

where S_W and S_B are the within-cluster and between-cluster scatter matrices, respectively. Being the ratio of the spread between classes relative to the spread within each class, the measure J increases monotonically as classes become increasingly more separable.

With this approach, the authors were able to demonstrate that cycling each sensing film through the 32 temperatures shown in **Figure 4A** did not necessarily create information that spanned 32 different dimensions. The responses were highly correlated and information seemed to be grouped based on temperature ranges; all lower temperature responses of a film type provided similar information that differs from that available from high temperature signals. On the other hand, cross-correlations computed across materials were comparatively lower than self-correlations.

Therefore, taken together with the results from the dimensionality reduction analysis, these results suggested that different materials provide orthogonal information about the target analytes.

There is a last important remark to emphasize here. The form of dependence of the objective function of the methods presented here on their respective optimizable control parameters is initially unknown, yet to be inferred from a provided training set containing labeled measurements from the same sensor array at representative values of the parameter itself. As a consequence, any change in the problem setup, e.g., addition or removal of an analyte class, necessitates a re-calculation of the control parameter with the updated dataset, meaning that the sensor array must be re-conditioned for each class configuration. This outcome, nonetheless, is normal for any optimization solution considered, meaning that the found solution is to be customized to the set of analytes being analyzed. In any case, the statistical methods presented in this section provided a generalizable methodology for designing and evaluating array-based solutions for a wide variety of specific detection problems. Ultimately, it is envisage that the advances generated by these methods are critical to the production of pre-programmed micro-sensors for non-invasive, real time, multi-species recognition relevant to homeland security, and other applications involving trace analyte detection in complex chemical cocktails.

ON THE OPTIMIZATION OF SINGLE CHEMICAL SENSORS

Much less attention has been paid to the optimization of metal-oxide sensors as a single device. As stated at the beginning of this review, and more specifically in Section "Metal-oxide Gas Sensors and Temperature Dependence," there are an extensive number of articles reporting empirical studies dealing with dynamic features obtained from transient responses, e.g., different temperature waveforms patterns and stimulus frequencies implemented as a countermeasure to the effects of selectivity and reproducibility encountered in gas sensors. There is no doubt that a high variance in response is detrimental in most chemo-transduction applications that must be tackled. This general treatment, though, constitutes only one facet of the sensor optimization problem that does not necessarily yield better performance in the odor identification task. The reason for this is that a reduction in the response variance does not ensure a non-overlapping class configuration in the feature space. Therefore, to maximize classification performance, one needs a more comprehensive formulation that quantifies the separation of specific odor classes in the sensor response. Both of these aspects have been covered in literature under the notion of "optimization." The thematic issues relevant to these works will be reviewed here.

OPTIMIZATION OF EXCITATION PROFILES

Chemical sensing can benefit from a variable-temperature signal generation. In most of the cases, the temperature variation (a.k.a. temperature modulation) has been approached empirically by implementing various temperatures waveforms and stimulus frequencies (Sears et al., 1989a,b, 1990; Nakata et al., 1992, 1996, 1998; Semancik and Cavicchi, 1999). Although the results achieved by such a technique are very promising, in most of the reported works the selection of waveforms and the frequencies used to modulate sensor temperature has been conducted in a non-systematic way (Davis, 1991; Corcoran et al., 1998; Choi et al., 2002; Fort et al., 2002, 2003; Huang et al., 2003). Even the selection of features from the sensor transients is a somewhat obscure process. Therefore, since these selections are based on a trial and error procedure, there is no way to ensure that the modulation frequencies, modulation depth, or features chosen are the optimal for a given application. Very few authors, though, have systematically addressed this problem by suggesting different optimization strategies.

The first approach to review, in this context, is the one implemented by Kunt et al. (1998) and Cavicchi et al. (1996), who pioneered the development of an optimization method for temperature-modulated micro-hotplate gas-sensor devices. They implemented a two-step optimization process for determining the optimal temperature trajectory (or trajectories) that would exploit the information characteristics contained in the sensor dynamics when operated in temperature-pulsed mode. In particular, the authors sought to optimize the said sequence by adapting the pulse amplitude, pulse duration, delay between two consecutive pulses, and number of pulses in a cycle to better discriminate between two different gaseous analytes: ethanol and methanol. In the first stage, the authors introduced a black-box dynamical model of the sensor from input–output experimental data; they input a temperature programmed excitatory signal to the heating element of a single micro-hotplate sensor device and collected the sensor conductance in presence to vapors of methanol and ethanol. The authors then sought to predict the next conductance value of the sensor response (y_{i+1}) from the previous values of the conductance $\left\{y_k\right\}_{k=i}^{i-n_y+1}$, as well as the next and previous values of the temperature set points $\{u_k\}_{k=i+1}^{i-n_u+2}$, thus,

$$y_{i+1} = F\left(y_i, y_{i-1}, \ldots, y_{i-n_y+1}, u_{i+1}, u_i, \ldots, u_{i-n_u+2}\right), \tag{11}$$

where n_y and n_u represent the model order in the input and output, respectively.

Utilizing different dynamic modeling methods, a suitable model $F(\cdot)$ was then built from the experimental data collected, being the wavelet network (WNET) method, for this particular case, the most accurate of all the tested methods. This initial model was further trained to set the final parameter values of the model $F(\cdot)$ and then used to simulate the sensor response to different temperature programs. In the second stage, the authors then implemented an off-line optimization routine to find the "optimal" temperature profile $\{u_i\}_{i=1}^{T}$ that maximizes the distance between the (simulated) temperature-modulated sensor responses to the targeted gases, thus

$$\{u_i\}_{i=1}^{T} = \underset{u_1,u_2,\ldots,u_T}{\arg\max}\ d\left(y^{\mathrm{MeOH}}, y^{\mathrm{EtOH}}\right), \tag{12}$$

where y^{MeOH} and y^{EtOH} are the conductance responses predicted by the WNET models for methanol and ethanol, respectively.

Above, the search space for this optimal temperature profile is over a limited subset of realizable temperature pulses (e.g., lower and upper limits are chosen based on the sensor structure) and

under the constraint that two consecutive pulses cannot differ in more than 40°C in order to avoid drastic changes in the surface. Finally, the metric used to quantify the said distance is the normalized sum of squared differences (NSSD) between the two response curves

$$\mathrm{NSSD} = \sum_{i=1}^{n} \frac{\left(y_i^{\mathrm{MeOH}} - y_i^{\mathrm{EtOH}}\right)^2}{n}, \tag{13}$$

where n is the number of temperature pulses in a cycle.

Figure 5 shows the optimal temperature profile to discriminate ethanol and methanol that was computed and validated through experimental measurements. This temperature profile (see the temperature profile shown in **Figures 5C,D**) produces methanol and ethanol responses that are out of phase (i.e., easy to discriminate). On the other hand, when applying a simple linear ramp (see **Figures 5A,B**), the sensor responses to ethanol and methanol were highly overlapped (i.e., becoming a non-trivial case of discrimination).

Although this methodology is systematic and should be applicable to other analytes, its application to the qualitative and quantitative analysis of multi-component mixtures is not straightforward. The fact that the method relies on the construction of good predictive response models, complicates the optimization process for multi-gas, concentration variant environments.

With this motivation, more recently Vergara et al. (2005a,b, 2007a, 2008) introduced, in a series of works, a system-identification method for optimizing the temperature-modulation frequencies in order to solve a given gas analysis problem. The optimization method consisted of utilizing one of the most useful types of periodic signal for process identification, the pseudo-random sequences of maximum length (PRS-ML), either binary or multi-level, to determine the most suitable temperature-modulation frequencies for discriminating and/or quantifying a number of specific target compounds at different concentrations.

Pseudo-random sequence signals are the most popular choice for the persistently exciting perturbation signals required in system identification. The most common application of these sequences is the identification of linear systems. In particular, the pseudo-random signal sequences considered by the authors in their optimization scheme are based on maximum-length q-sequences, either binary or multi-level[1], the generation and properties of which were described by Zierler (1959). The relevant theory behind these signals is based on the algebra of finite fields. When q (the number of levels) is a prime, the digits of the sequence are the integers 0, 1, ..., $(q-1)$ and the sequence can be generated by a q-level, n-stage shift register with feedback to the first stage consisting of the modulo q sum of the outputs of the other stages multiplied by coefficients a_1, ..., a_n, which are also the integers 0, 1, ..., $(q-1)$. The length (or period) of a maximum-length sequence is $q^n - 1$, which signifies that the sequence repeats itself after $q^n - 1$ logic values. The pseudo-random sequences are periodic, deterministic signals that have a flat power spectrum over a large frequency range. These properties imply that a PRS-ML shares some properties with white noise, but with the advantage of being repeatable, which makes these signals even more attractive and suitable for the system-identification task ahead. The generator of such a sequence and an example of a 5-level sequence (fragment) are shown in **Figures 6A,B**, respectively. Notice how the initial state of the shift register can be any combination of length n of the values 0, 1, ..., $(q-1)$, with an exception made of $n =$ zeros, and that each combination of these values appears as the state of the register exactly once during a period of the PRS-ML (Godfrey, 1993).

The impulse response, $h(t)$, is the main descriptor of a linear invariant system. Among the different strategies to estimate the impulse response, noise based methods allow to excite the system under study for enough time to supply it with the necessary energy to obtain a good estimate of $h(t)$. By using white noise as excitation signals, one ensures that there is a homogeneous distribution of the energy over a large frequency range. Since PRS-ML signals have a low crest factor (i.e., low peak-to-average factor), they minimize the risk of saturating the system under study, which, in practice, means that these signals contain energy enough to obtain a good signal-to-noise ratio in a wide frequency range (i.e., measurement with high dynamic range) and that they avoid possible sensor non-linearities caused by signals with high crest factors (e.g., impulsive signals). Therefore, since these excitatory noise signals are deterministic, reproducible results are expected to be obtained, provided that the conditions of the system under analysis remain unchanged.

The power spectrum envelope of a PRS-ML is almost flat up to a frequency equal to $0.45 \times f_c$, where f_c is the frequency of the clock signal applied to the shift register used to generate the signal. The power spectrum is discrete and the separation between spectral lines (i.e., the spectral resolution) is f_c/L, where L is the length of the PRS-ML. **Figure 6C** shows the power spectrum of a PRS-ML signal, where, as observed, the power spectrum envelope is similar to the power spectrum of white noise up to the -3 dB cut-off frequency, which in this particular case is equal to $0.45 \times f_c$.

When the pseudo-random sequence is a maximum-length signal, the impulse response estimate, $\hat{h}(n)$, can be obtained by computing the circular cross-correlation between the excitatory signal, $x(n)$, and the response signal, $y(n)$. The circular cross-correlation of two sequences x and y in $\Re^L$ may be defined as,

$$\hat{h}(n) = \frac{1}{L} \sum_{l=0}^{L-1} y(l+n)x(l),\ n = 0,\ 1,\ 2, \ldots,\ L-1. \tag{14}$$

Above, the cross-correlation is circular since $l+n$ is interpreted as modulo of L, where L is the length of the sequence and that can be optionally utilized as a normalization factor. The circular cross-correlation between the input and output sequences can readily be interpreted in terms of $\hat{h}(n)$, since the autocorrelation function of the PRS-ML signal is of approximately impulsive form.

[1]One of the main reasons for considering signals with more than 2 levels is that multi-level signals provide the possibility of identifying a better estimate of the linear dynamics of a process with non-linearity than the binary sequences and that they can also be of use in the identification of the non-linear characteristics themselves (Vergara et al., 2005b, 2007a).

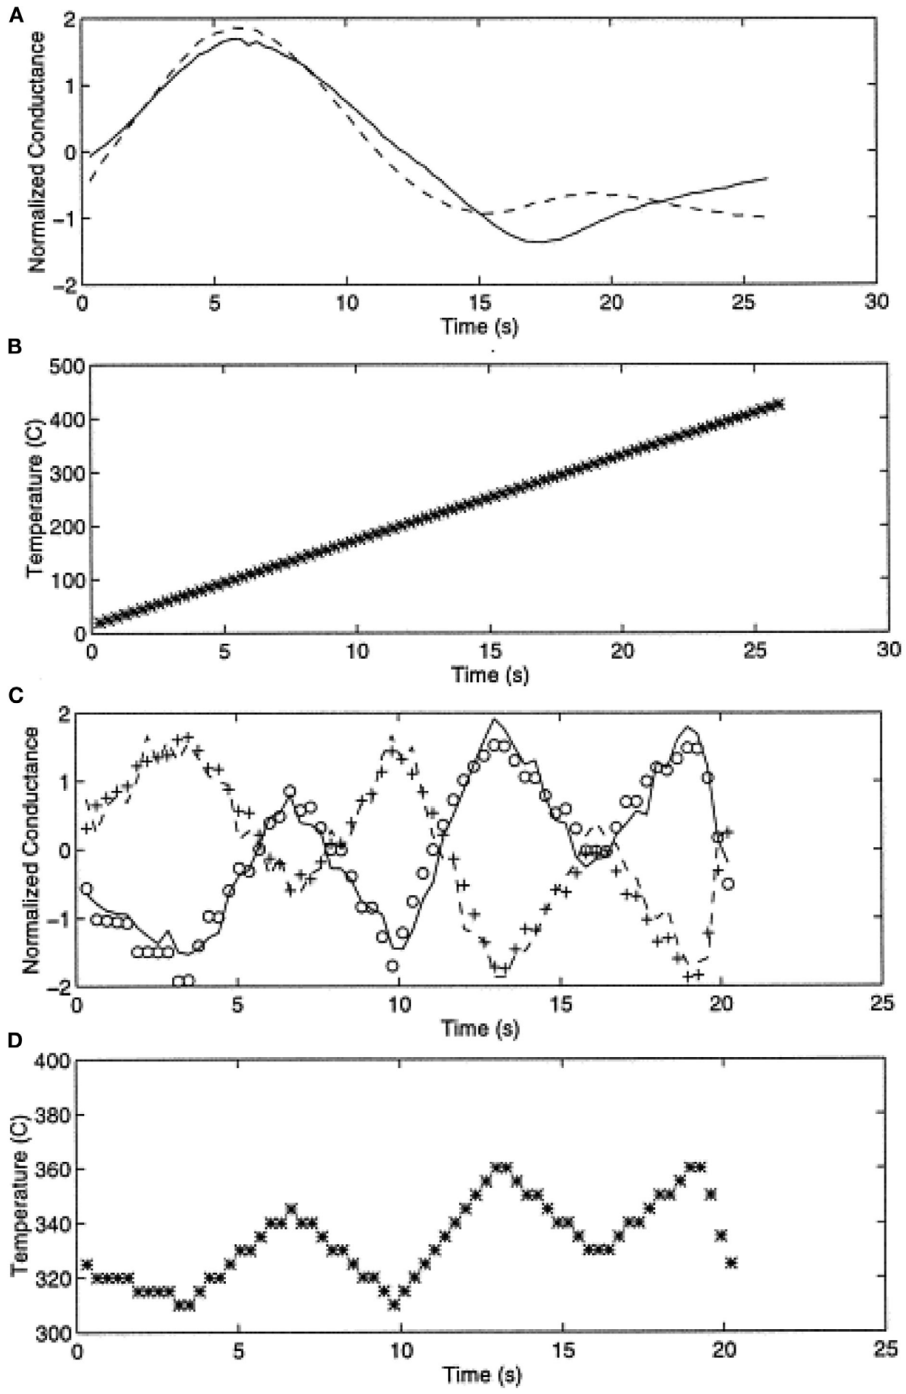

FIGURE 5 | (A) Normalized conductance response to methanol (solid) and ethanol (dashed) upon applying a linear temperature ramp as shown in **(B)**. **(C)** Actual experiments with methanol (solid line) and ethanol (dashed line) gases, model predictions are shown by circles for methanol and plus for ethanol models; **(D)** the optimum temperature profile derived from the off-line optimization process. Note the dramatic improvement in discrimination between **(A)** and **(C)**. Figure reprinted from Kunt et al. (1998), Copyright 1998, with permission Elsevier Science.

The whole optimization method proposed is illustrated in **Figure 7**. In a practical instance, it works as follows. First, a voltage PRS-ML signal is applied to the heating element of a micro-hotplate gas sensor while the sensors are exposed to various target compounds (e.g., nitrogen dioxide, ammonia, ethylene, ethanol, acetaldehyde, and their binary mixtures), hence ensuring

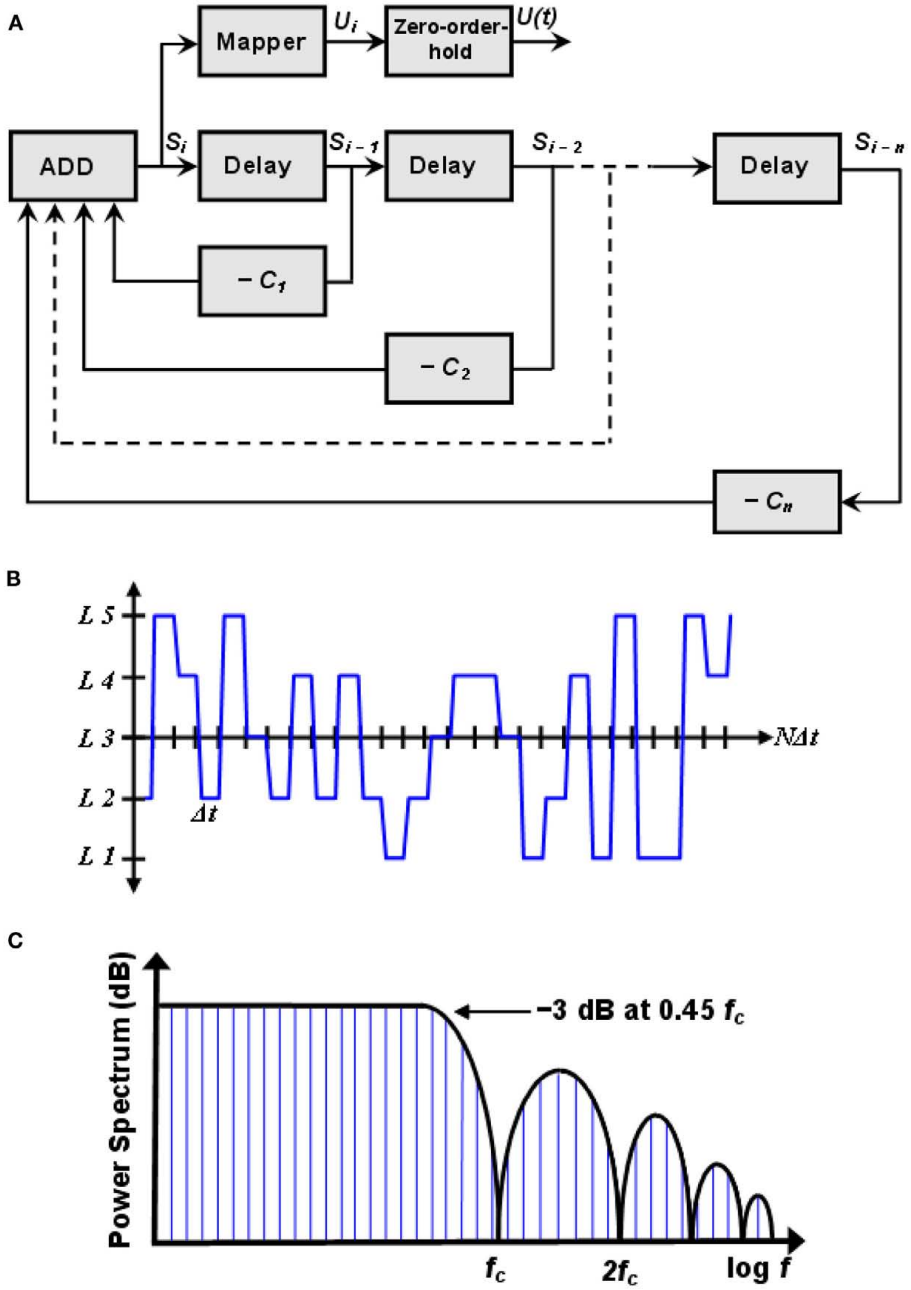

FIGURE 6 | (A) A *q*-level pseudo-random maximum-length sequence generator. **(B)** Fragment of a 5-level pseudo-random sequence. **(C)** Discrete power spectrum of a PRS-ML signal. Spectral resolution is f_c/L, where f_c is the frequency of the clock signal applied to the shift register and L is the length of the sequence. Figure reprinted from Vergara et al. (2007a), Copyright 2007, with permission Elsevier Science.

that the sensor's working temperature is modulated over the whole wide frequency range considered. Then, for each individual target compound the impulse response $h(t)$ is computed as the circular cross-correlation between the excitation signal (PRS-ML) and the sensor response. Afterward, the absolute values of the FFT of the impulse response estimate are calculated, determining, in essence, which spectral components contain important information for the identification and quantification of gases. Finally, each individual frequency is ranked on the basis of its information content (between-class to within-class scatter ratio), and a subset of the most informative frequencies is selected. As the authors have shown in their results, this procedure succeeded in maximizing the discrimination and quantification of various gases and their mixtures using even a single sensor with its optimized set of modulating frequencies.

In addition, the authors extended their optimization study on a completely two-stage validation procedure, demonstrating thus the consistency and robustness of the method itself

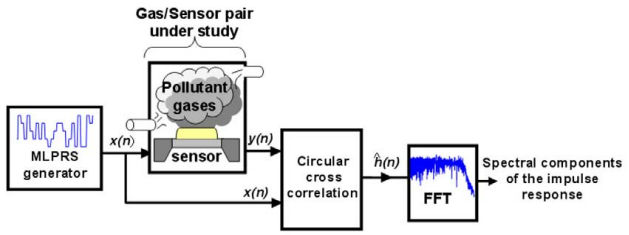

FIGURE 7 | Study of the sensor/gas system using MLPRS signals. The MLPRS voltage signal, $x(n)$, is input to the heating element of a micro-hotplate gas sensor. The transient of the sensor conductance (i.e., the response in the presence of gases), $y(n)$, is recorded. An estimate of the impulse response, $\hat{h}(n)$, can be computed via the circular cross-correlation of $x(n)$ and $y(n)$. Finally, by performing the FFT of $\hat{h}(n)$, the spectral components of the impulse response estimate are found. Figure reprinted from Vergara et al. (2007a,c), Copyright 2007, with permission Elsevier Science.

(Vergara et al., 2007c, 2008). The two-stage method consisted of, first, running the whole optimization process presented above utilizing a set of measurements pre-recorded with one micro-sensor array, and then, validate the resulting outcome employing a new set of measurements collected with a different sensor array with the same characteristics of the one used in the first stage. This new set of measurements, though, was based on the multi-sinusoidal temperature-modulating signal showed in **Figure 8A**, the frequencies of which were the reduced set of the optimal ones estimated in the first stage. The **Figures 8B,C** show the FFT spectra of the transient response of a sensor in the presence of two different analytes [acetaldehyde (50 ppm) and ethylene (50 ppm)]. Notice how the peaks in these plots correspond to the "optimal" temperature-modulating frequencies selected. As the variation of peaks' height (i.e., the pattern) between **Figures 8B** and **8C** shows, significant improvements in the classification and quantification capability of a single gas-sensor operated under the temperature-modulation scheme is attained. Particularly, their results revealed a classification performance of up to 98.2% even when a single sensor was used, and a shift of the odor concentration prediction of down to 0.92 ppm for single species and 2.81 ppm for binary mixtures. These illustrated performance improvements were expected, though, in the sense that, as the **Figures 8B,C** shows, a different pattern develops when different odorant species are measured, whilst the resulting pattern is preserved, to a large extent, when the vapor concentrations change, meaning that the illustrated evaluation is a reliable indicator of the improvement of the classification capability of the sensors when their operating temperature is modulated at the frequencies selected.

As a final remark to emphasize here, it was demonstrated that for each gas-sensor pair, the modulating frequencies selected are related to the characterization of the interaction between the metal-oxide layer and the gas, e.g., film microstructure; surface diffusion; and reaction kinetics. Even though the method was implemented for the analysis of the specific qualitative and quantitative task earlier described, this optimization procedure is generic and could be applied to many qualitative and quantitative gas analysis applications.

OPTIMIZATION OF THE OPERATING TEMPERATURE: INTERNALLY TUNING THE CHEMICAL SENSORS

It has long been known that varying or setting different values of the sensor's operating temperature affects all the aspects of the sensor response, including its selectivity and sensitivity to different volatile compounds (i.e., the sensor's ability to encode the odor information), as well as its reproducibility. For example, carbon monoxide (CO) is usually best detected at lower operation temperatures (e.g., 250°C) when using a tin dioxide based sensitive layer, whereas higher temperatures (e.g., 350°C) are used for monitoring hydrocarbons such as methane among others. In view of this, different strategies, such as the idea of periodically changing the sensor working temperature, have been implemented to maximize the performance of the sensors. However, despite the promising results obtained in all these previously cited attempts, one question remains unanswered: given a metal-oxide based chemical sensor, how does one select the best (i.e., the optimal) operating temperature (or temperatures) for fast and reliable discrimination or quantification of chemical species? One conceivable manner to address this issue is to empirically vary the operating temperature through all the possible values available in the sensor so that its response to each gas is maximized (Cavicchi et al., 1996; Maziarz and Pisarkiewicz, 2008). However, this may be an expensive and inefficient solution because it does not guarantee the improvement of the performance of the sensors.

Undoubtedly, heightened sensitivity to a spectrum of chemical hazards is necessary for the detection of analytes at relevant concentrations. However, this general treatment constitutes only one facet of the problem, a substantial selectivity is also necessary to rapidly and accurately perform the odor identity representation task. The reason for this is that an increase in the response does not ensure a non-overlapping class configuration in the feature space. Therefore, to maximize the classification performance, one needs a more comprehensive formulation that quantifies the separation of specific odor classes in the sensor response.

Following this scheme, Vergara et al. (2009b, 2010) formulated an optimization method to select, for a single sensor, the best operating temperature to discriminate a given set of odorants. The authors presented a rigorous way of selecting the best operating temperature for a chemical sensor. The method hinges on an information measure widely used in information theory, namely the relative entropy or Kullback–Leibler divergence (KL-divergence; Kullback and Leibler, 1951), a measure index that rates the difference between two probability distributions. Since these probability distributions may belong to one of the disjoint classes of interest in a particular odor universe, this annotated quantitative measure shows how odors are encoded in every odorant chemo-receptor and how distinguishable they are from each other at different parameter values. Tuning a control parameter, such as the sensor's operating temperature, will maximize such a difference, yielding thereby a substantial improvement in the classification performance (separation of classes) and reproducibility of the process. In particular, using a metal-oxide gas sensor in an odor-discrimination instance, the authors demonstrated the proposed criterion by studying the impact of adjusting the sensing parameter on the odor-sensor pair interaction and on the confidence of the information yielded by the sensor individually.

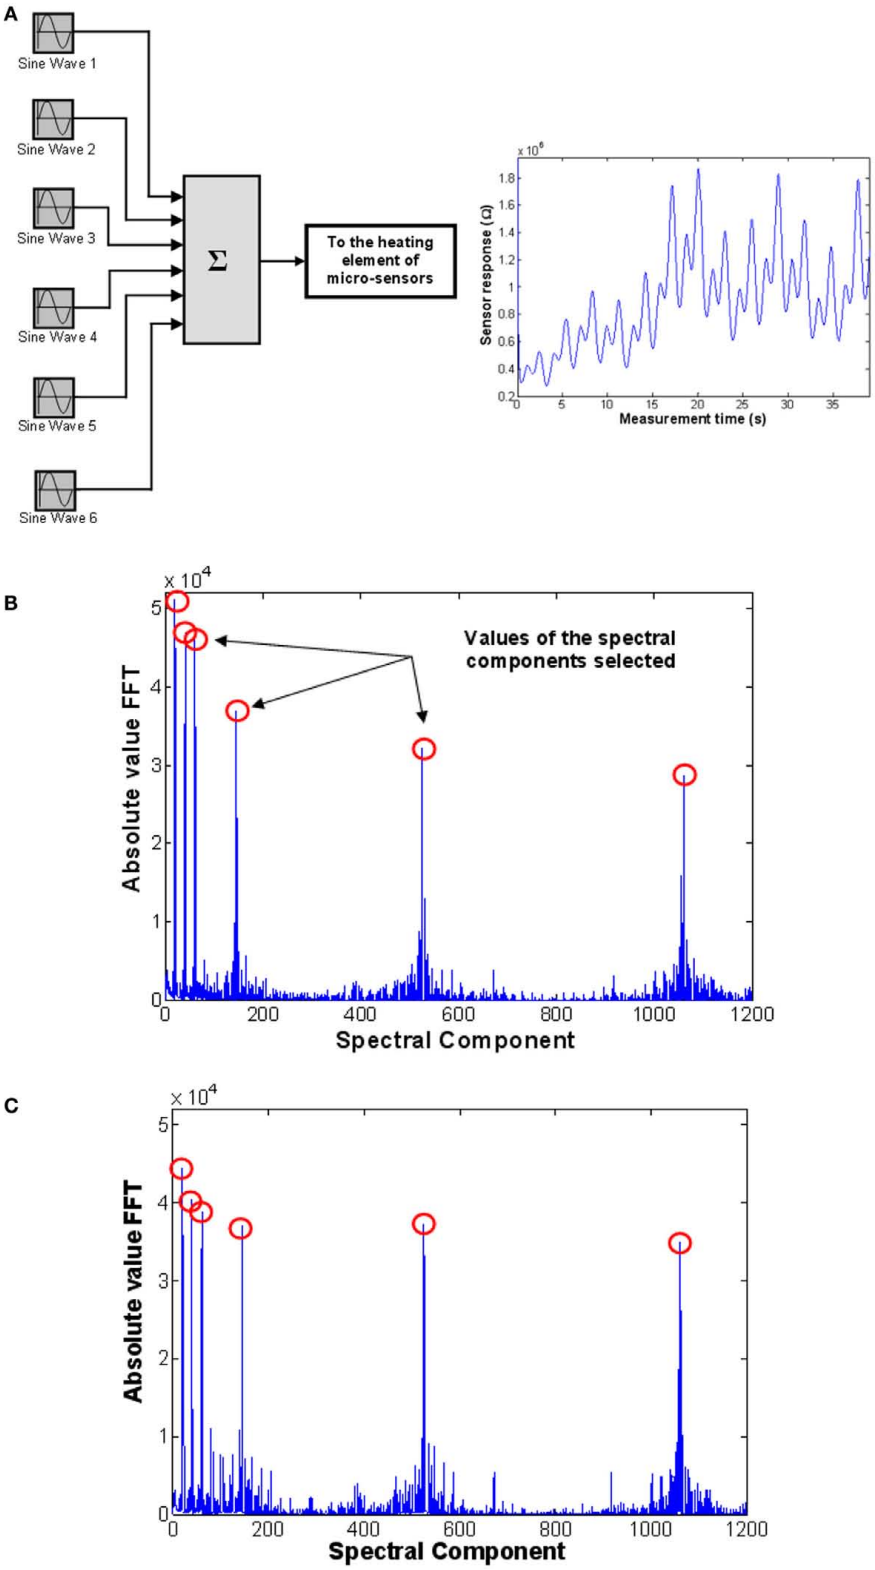

FIGURE 8 | (A) Setup used to generate a multi-sinusoidal signal, which consists of the sum of six sinusoids of identical amplitudes and different frequencies. The signal is applied to the heating element of the micro-sensors studied. FFT (absolute value) of the transient response of a temperature-modulated WO_3 micro-hotplate sensor in the presence of **(B)** 50 ppm acetaldehyde and **(C)** 50 ppm ethylene. Figure reprinted from Vergara et al. (2007a, 2008), Copyright 2007 and 87, Copyright 2008, with permission Elsevier Science.

The KL-divergence is a very well-known index for class separation that is a non-commutative measure (a "distance" in a heuristic sense) of the difference between two distributions: a conceptual reality [probability distribution $g(\cdot)$] and an approximate model [probability distribution $h(\cdot)$]. For two continuous functions qualifying as probability distributions, the KL-divergence is defined by the integral:

$$\mathrm{KL}(g||h) = \int_{-\infty}^{\infty} g(x) \log \frac{g(x)}{h(x)} dx, \tag{15}$$

where KL is the measure of "information" lost when a model $h(\cdot)$ is used to approximate reality [i.e., model $g(\cdot)$; Kullback and Leibler, 1951].

The utility of the KL-divergence is based on a certain number of properties that make it unique for measuring the difference between two probability distributions. For example, this approach can account for a number of key characteristics of a response, including, for example, higher order moments (e.g., skewness) or multi-modality, which in turn may be involved in the response distribution (at least in odor representation), causing thus loss of information. However, the measure is still not commutative, i.e., $\mathrm{KL}(g(x) \,||\, h(x))$ is in general different from $\mathrm{KL}(h(x) \,||\, g(x))$; therefore, the KL-divergence is not a legitimate metric by itself. As a consequence, a symmetrized version, namely the KL-distance, can be readily composed after a straightforward manipulation given by:

$$\mathrm{KL}(g, h) = \frac{1}{2}\mathrm{KL}(g||h) + \frac{1}{2}\mathrm{KL}(h||g), \tag{16}$$

which the authors adopted as a measure of the class-conditional distributions' separation for the specific purpose.

In a context C (i.e., the likelihood of occurrence of each odor stimulus from a finite list of analytes whose classes are known) in which one is trying to discriminate two compounds that is complicated by the similarities/overlaps among the class-conditional distributions, the KL-distance index, given in Eq. 16, constitutes an accurate measure of discrimination, hence a good indicator of the classification performance for any unbiased classifier. Therefore, given the simplest two-dimension discrimination problem (i.e., a two-class discrimination task), when the class-conditional distributions depend on a measurement parameter (e.g., operating temperature in metal-oxide gas sensors), maximizing the KL-distance is a valid objective function for tuning that parameter (see **Figure 9**).

Using a binary classification instance as a case of study may be very convenient from many perspectives; in odor representation, however, this assumption may be very unrealistic. When the number of classes (i.e., the possible outcomes of the identification problem) is more than two, the KL-distance should be generalized to promote the dispersion of the whole classes. The authors have addressed this issue by replacing (16) with the sum of pair-wise distances, thus:

$$\mathrm{CKL} = \sum_{i,j=1}^{|C|} \mathrm{KL}(g_i, h_j), \tag{17}$$

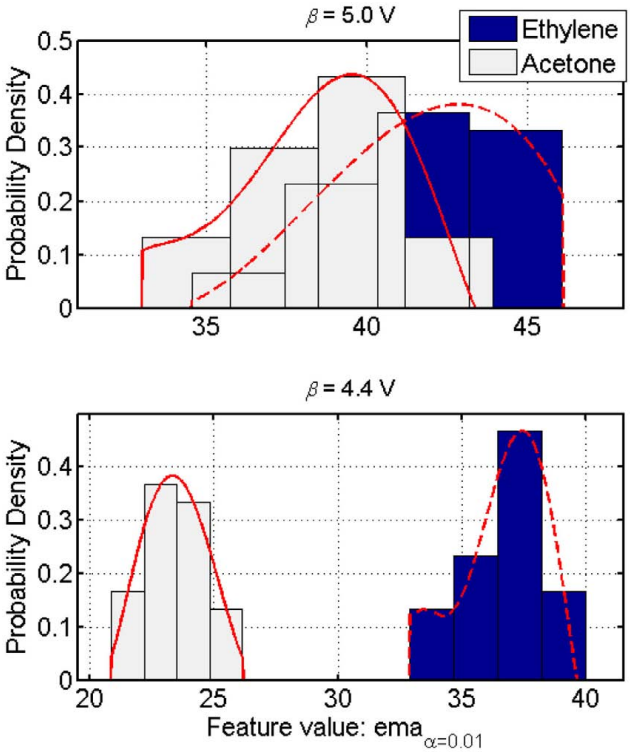

FIGURE 9 | Class-conditional probability distributions in a two-class discrimination instance. Each class response models the histograms by a normalized fifth-order polynomial (plain and dashed lines). These models accurately approximate the sensor's response to an odor class while accounting for the asymmetry (i.e., skewness) in the distribution. The KL-distance index then captures the influence of the operating parameter on the separability of such distributions. Maximizing this index for a pair of distributions results in a better discrimination between the corresponding two classes (top versus bottom figures). Figure reprinted from Vergara et al. (2010), Copyright 2010, with permission Elsevier Science.

where $|C|$ denotes the number of classes in the problem and $g_i(\cdot)$ and $h(\cdot)$, $i \in C$, $j \in C$, are the class-conditional distributions, each potentially depending on the operating parameter (e.g., the sensor's operating temperature). The CKL quantifies the difficulty of the classification problem. When this quantity is large, an arbitrary classifier is expected to perform with higher accuracy, since, relative to a small CKL, the distribution within each class shrinks whilst the distance among the considered classes increases in the feature space.

Assuming that β denotes an intrinsic parameter of a sensor device that alters the response characteristics (see **Figure 9**), then the problem configuration CKL is expected to be sensitive toβ itself. Hence, the value

$$\beta^* = \arg\max_{\beta} \mathrm{CKL}(\beta), \tag{18}$$

defines an optimum operating condition for the classification problem at hand.

In demonstrating their optimization scheme, the authors applied the criterion (18) to optimize the operating conditions

of commercialized metal-oxide gas sensors (sensors provided by Figaro Engineering Inc., Japan, http://www.figaro.co.jp). In particular, they have examined the performance of each gas sensor in a six-class classification problem, comprised by six different analytes dosed at different concentrations (i.e., ethylene, ethanol, and toluene dosed at 10 ppm; acetone and acetaldehyde at 100 ppm; and ammonia at 120 ppm). They thus studied the impact of adjusting its sensing parameter β on the odor-sensor pair interaction and on the confidence of the information yielded by the sensor individually.

In principle almost any controllable variable that alter or modify the operating characteristics of the sensor, such as the environment temperature, flow rate, or even construction methodologies, can be used as a parameter of the response profile that can be tuned to improve the processing performance. However, in this popular odor sensing technology, it is very well-known, and proved in many empirical works (see, e.g., works from Sears and Nakata), that there is a strict dependence of the sensor response on its operating temperature (temperature normally ranging in high orders of magnitude, e.g., 400°C, responsible of the adsorption/desorption reaction occurring at the micro-porous surface of the sensor in response to an analyte). Accordingly, having such an easy way of interacting with the sensor, the most natural way to optimize the sensor device is with respect to this parameter (i.e., the sensor's operating temperature), assuming that all the other parameters remain constant. Since the sensor packaging does not permit direct access to this temperature, its tuning can be achieved via a resistive heater element with controllable voltage, which has a deterministic one-to-one mapping with the actual active layer temperature[2]. Accordingly, the authors have considered this heater voltage and the operating temperature interchangeably as the sensing conditioning parameter β to be optimized.

To demonstrate the optimization scheme in a practical instance, the authors established the following procedure: Initially, the form of the dependence of CKL on β is initially unknown, yet to be inferred from a provided training set (containing labeled measurements) from the same sensor at representative β values. For each sensor the authors then compiled a comprehensive dataset containing the analytes described above. Each set of time series contained 30 independent measurements taken from each class at each of the 13 sensor operating temperatures corresponding to the heater voltages $\beta \in \{3.8, 4.0, \ldots, 6.2\,\mathrm{V}\}$[3]. Thus, the authors represented each of the chemo-sensory records, associated with each odor class and operating temperature, as independent and identically distributed (i.i.d.) samples, from which the class-conditional distribution is derived. Then, they modeled each odorant class by a polynomial fit to the histogram of previously collected samples from that odorant type. In particular, they consider a fifth-order polynomial to represent odorant class/histogram relation. Finally, by plugging these functions into Eq. 17, the CKL criterion (18) was implemented, and the maximum β value obtained, yielding thus the optimal operating temperature values for the particular discrimination task. As a measure of separability, the evaluation of the CKL for the six classes at each temperature β yielded the profile shown in **Figure 10** (dashed lines). Based on this evaluation, the best operating condition for each sensor to distinguish between the set of classes is determined by the maximum value of their respective curves.

[2]A look-up table with a deterministic one-to-one mapping of the actual active layer temperature and the controllable heater voltage was provided in Figaro Engineering Inc., Japan, http://www.figaro.co.jp

[3]The operating temperature values selected perform a dataset with a temperature resolution (i.e., separation value among the temperatures evaluated) of 20°C, under the assumption that this temperature resolution suffices in making a reliable inference of the dependence of the CKL on the parameter value β.

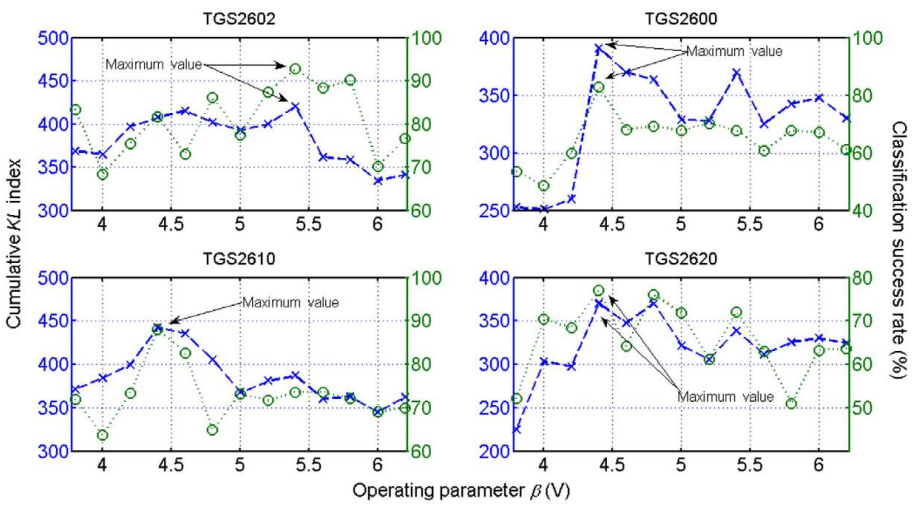

FIGURE 10 | Observed discrimination performance of a linear-SVM classifier of each sensor on the six-class identification problem (dotted lines). Profiles for each sensor as estimated by the CKL index (dashed lines) with respect to β. Based on the proposed criterion, the optimal operating condition that best discriminates between the set of classes can be determined for each sensor individually by obtaining the maximum value of the curve shown. Figure reprinted from Vergara et al. (2010), Copyright 2010, with permission Elsevier Science.

To demonstrate the consistency and robustness of the optimization method, the authors conducted a validation process that consisted of measuring the correlation between the information given by the optimization criterion (18) and the performance given by an arbitrary linear support-vector classifier based classifier. As observed in **Figure 10**, the results yielded by the SVM classifier (see **Figure 10**, dotted lines) follow a similar pattern to the estimated measure CKL index (**Figure 10**, dashed lines) in the sense that their extreme points occur at the same β values. The ordering of these points in magnitude was also preserved to a large extent, meaning that the proposed measurement is a reliable indicator of the classification performance at almost all temperatures within the range.

In addition, the authors validated the whole optimization process with the second dataset 4 months later. To validate the results, they re-calculated the proposed cost function CKL of the second dataset by applying the same procedure described above. Based on this re-evaluation, the best operating condition to distinguish between the new set of classes was determined, and compared to the performance yielded by the linear-SVM classifier. As can be seen in **Table 1**, the results obtained in the validation stage perfectly matches with the information given by the classifier, showing the consistency of the method. These results indicate that the proposed CKL measure-index is optimum for any complexity of probability distribution models; hence, for any type and number of odorants as well as any type of sensor technology with a conditioning parameter, provided that these models are accurate in identifying the response distribution. Nevertheless, the method could be extended to an arbitrary classification instance as needed, including complex odors at different concentrations or mixtures of gases, provided that a sufficiently representative database of relevant measurements is available. It is also important to emphasize that the solution β^* does not impose a particular classification method. Therefore, the parameter value resulting from the maximization of Eq. 18 simplifies the task of an arbitrary unbiased classifier.

It is important to comment on one last issue here. An operating condition is optimal for a well-defined task. If this task changes then the best condition should be re-calculated. This applies, nonetheless, to any optimization method, not just this one. For example, considering a generic classifier training instance, if the training data changes (e.g., some data turns out to be invalid or relabeled), then the device needs to be re-trained in order to determine the optimal performance. In this optimization case, a re-calculation of β^* with the updated dataset is therefore needed, too.

ACTIVE-SENSING OPTIMIZATION

The idea of applying sophisticated signal processing procedures and optimization strategies to ameliorate the performance of metal-oxide gas sensors has been around for more than two decades. Researchers have since used a wide array of dynamic features obtained from transient responses, but most of these studies have been empirical. To the best of our knowledge, very few studies have proposed systematic approaches to optimizing the sensor performance as a single device (Cavicchi et al., 1996; Kunt et al., 1998; Vergara et al., 2005a,b, 2007a,c, 2008, 2009a,b, 2010). These methods, though, require that the optimization be performed offline; therefore, they cannot adapt to changes in the environment. In view of this, a novel active-sensing approach that can optimize the temperature profile online (i.e., as the sensor collects data from its environment), has recently emerged in literature. The most relevant works on this thematic issue are reviewed in this section.

Active-sensing strategies are inspired by the fact that perception is not a passive process (Gibson, 1979), but an active one, in which an organism controls its sensory organs in order to extract behaviorally relevant information from the environment (see **Figure 11A**). Active sensing has been traditionally used in robotics and computer vision, in which the localization and navigation tasks, on the one hand, and the recognition of three-dimensional (3-D) objects from 2-D image, on the other hand, respectively, is a recurrent theme (Paletta and Pinz, 2000; Denzler and Brown, 2002; Floreano et al., 2004). In chemical sensing, however, it has received only minimal attention. In one of the earliest studies, Nakamoto et al. (1995) developed a method for active odor blending, where the goal was to reproduce an odor blend by creating a mixture from its individual components. The authors developed a control algorithm that adjusted the mixture ratio, so the response of a gas-sensor array to the mixture could matched the response to the odor blend.

It was not until 2010, when Gosangi and Gutierrez-Osuna (2009, 2010) proposed an active-sensing approach to optimize the temperature profile of metal oxide sensors in real time, as the sensor reacts to its environment. To see how their approach works let us consider the problem of classifying an unknown gas sample into one of M known categories $\{\omega^{(1)}, \omega^{(2)}, \ldots, \omega^{(M)}\}$ using a MOX sensor with D different operating temperatures $\{\rho^{(1)}, \rho^{(2)}, \ldots, \rho^{(D)}\}$. To solve this sensing problem, one typically measures the sensor's response at each of the D temperatures and then analyzes the complete feature vector $x = [\chi^1, \chi^2, \ldots, \chi^D]^T$ with a pattern-recognition algorithm. Although straightforward, this passive sensing approach is unlikely to be cost effective because only a fraction of the measurements are generally necessary to

Table 1 | Optimal operating parameter values β selected versus the observed classification performances for each metal-oxide gas sensor given by the linear-SVM classifier during the validation stage.

Sensor type →	TGS2602	TGS2600	TGS2610	TGS2620
Optimal parameter value β using the CKL-distance (V)	5.4	4.4	4.4	4.4
Validation discrimination performance rate in percent (CKL-distance)	90.50	84.55	87.68	94.23

The performance of the linear-SVM classifier was quantified for each sensor on each optimized parameter value. Table adapted from Vergara et al. (2010), with permission from Elsevier Science.

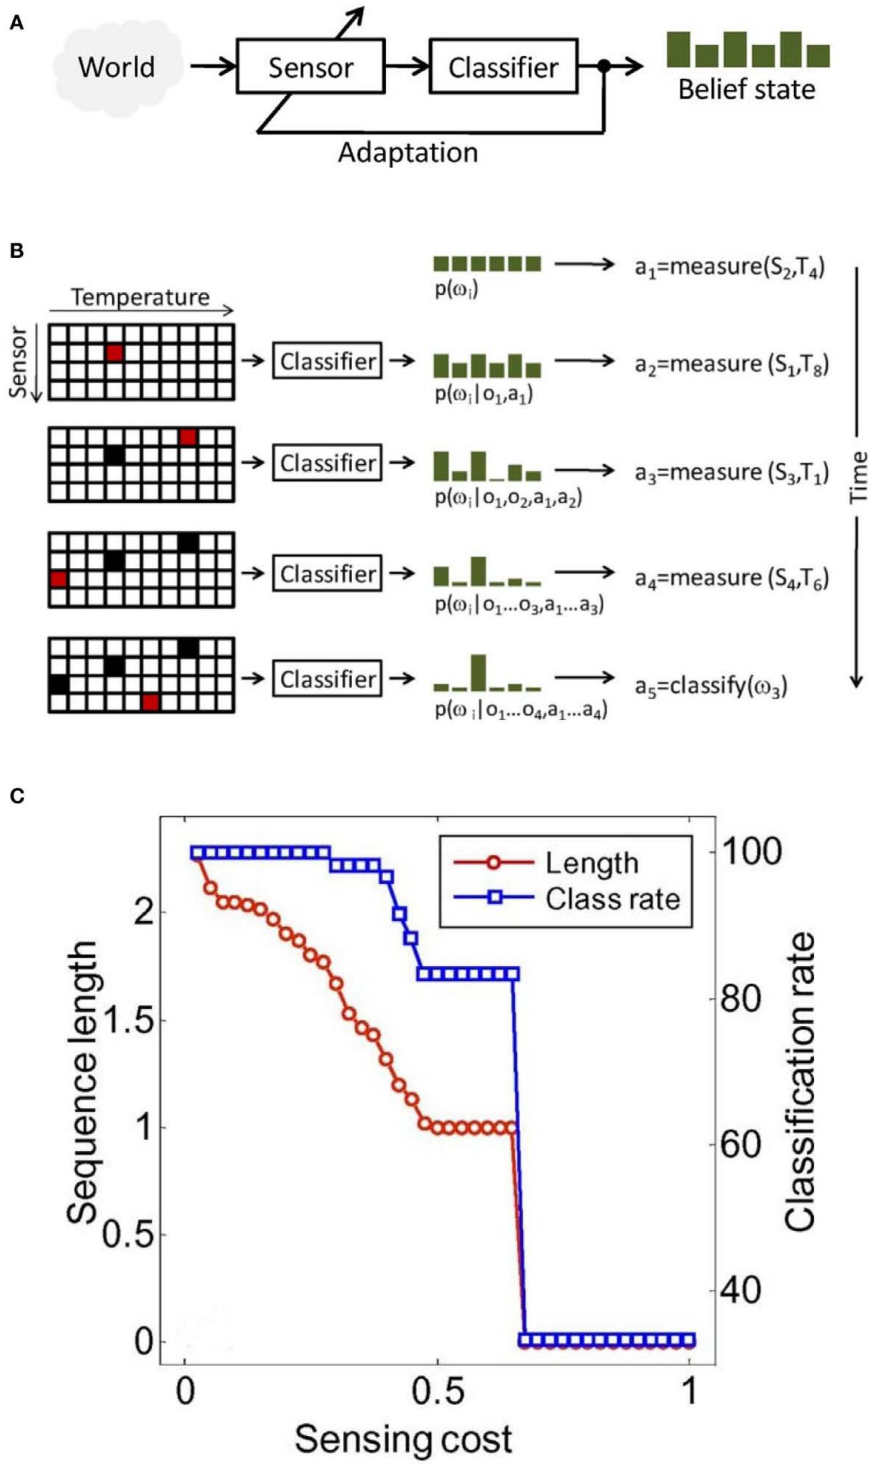

FIGURE 11 | (A) In active sensing, the system adapts its sensing parameters based on its belief about the world (e.g., class membership of a stimulus). **(B)** Illustration of active-classification with an array of four metal-oxide gas sensors, 10 temperatures per sensor, and a discrimination problem with six chemicals. At time zero, no information is available except that classes are *a priori* equiprobable: $p(\omega^{(j)} = 1/6)$. Based on this information, the active classifier decides to measure the response of sensor S2 at temperature T4, which leads to observation o_1 and an updated posterior $p(\omega^{(j)}|o_1,a_1)$. After four sensing actions, evidence accumulated in the posterior $p(\omega^{(j)}|o_1 \ldots o_1, a_1 \ldots a_1)$ and the cost of additional measurements are sufficient for the algorithm to assign the unknown sample to class ω (3). In this toy example, accurate classification is reached using only 10% of all sensor configurations. **(C)** Classification performance and average sequence length as a function of feature acquisition costs. Figure adapted from Gosangi and Gutierrez-Osuna (2009) ©IEEE.

classify the chemical sample. Instead, the authors seek to determine an optimal sequence of actions $a = [a_1, a_2, \ldots, a_T]$, where each action corresponds to setting the sensor to one of the D possible temperatures or terminating the process by assigning the sample to one of the M chemical classes. More importantly, they seek to select this sequence of actions dynamically on the basis of accumulating evidence. This process is illustrated in **Figure 11B**.

In demonstrating their approach, the authors first model the sensor's steady-state response at temperature ρ_i to chemical $\omega^{(c)}$ with a Gaussian mixture:

$$p\left(x_i\omega^{(c)}\right) = \sum_{m_i=1}^{M_i} \alpha_{i,m_i}^{(c)} N\left(x_i|\mu_{i,m_i}^{(c)}, \sum_{i,m_i}^{(c)}\right) \tag{19}$$

where M_i is the number of Gaussians, and $\alpha_{i,m_i}^{(c)}, \mu_{i,m_i}^{(c)}, \sum_{i,m_i}^{(c)}$ are the mixing coefficient, mean vector, and covariance matrix of each Gaussian for class $\omega^{(c)}$, respectively. Given a sequence of actions $[a_1, a_2, \ldots, a_T]$, the authors assumed that the sensor progresses through a series of states $s = [s_1, s_2, \ldots, s_T]$ to produce an observation sequence $o = [o_1, o_2, \ldots, o_T]$. Each state s_i represents a Gaussian in (19) and is therefore hidden. Following this step, they modeled the dynamic response of a sensor to a sequence of temperature pulses by means of an input–output hidden Markov model. This is a machine learning technique that can be used to learn a dynamic mapping between two data streams: (a) an input (temperature in this case) and (b) an output (sensor conductance). Once a dynamic sensor model has been learned, they then approach the temperature-optimization process as one of sequential decision-making steps under uncertainty, where the goal is to balance the cost of applying additional temperature pulses against the risk of classifying the chemical analyte on the basis of the available information. As a result, the problem is solved through a partially observable Markov decision process (Papadimitriou and Tsitsiklis, 1987).

Simulation results from this study are shown in **Figure 11C**; these results indicate that the method can balance sensing costs and classification accuracy: higher classification rates can be achieved by decreasing sensing costs, which in turns increases the length of the temperature sequence and the amount of information available to the classifier. One last think to comment on here is that the active-sensing approach proposed here has great potential in pioneering new strategies to be implemented in energy-aware chemical sensing networks using low-cost commercial sensors.

CONCLUSION AND OUTLOOK

Most of the work reviewed in this article has focused on optimization schemes for single sensors or sensor arrays within a single chemo-sensing system. Advances at the intersection of materials science, micro- and nano-technology, microelectronics, and chemical micro-sensory systems together with the explosion of mobile computing and the wireless communications capabilities will soon dramatically influence the deployment of massively distributed sensory micro-systems networks with capabilities that, only years ago, were confined to the lab bench of research laboratories. It can be concluded from the contents of this review that the integration of chemical micro-systems and wireless networks may, indeed, have a profound impact not only in different applications ranging from environmental monitoring to personal health care, but also in the formidable challenges present at the current practices of sensor management and data analysis. However, we believe that the sensor-based detection of chemical analytes in a dynamic real-world environment is a complex task, in which the interest of the final-user is to have reliable, user-friendly, and affordable sensory systems, irrespective of the internal system complexity. Therefore, the concept of an adaptive optimal sensory system can be most successful.

Adaptive sensory systems may be devices that include various chemical transducer types and sensor operation modes, the use of auxiliary sensors, and separation and pre-concentration units, which can respond or adapt their optimal operation to the occurring analysis situations or events. Then, in the instance where a certain odorant compound or a major interfering chemical is present, the sensor/feature selection, sensor operation mode, feature extraction, and data treatment should be able to adapt to this event so that their protocols execute in such a way that the best-possible target-analyte detection is achieved or that the influence of the interfering analyte can be recognized and minimized if not suppressed at all. In dealing with the just described issues, it may be very effective to purposefully select or deselect sensors, to find the optimal sensor's operation mode, or to use signal ratios or differential values instead of merely increasing the array size or the transducer diversity.

Another important observation to make from the optimization schemes reviewed here is that the chances for these optimization schemes to succeed, or in other words maximizing the performance in a gas-sensor identification system to be more precise, truly depends on multiple factors that all together determine the problem setup. These factors include the hierarchical level of the classification problem pursued (e.g., gas discrimination versus quantification and prediction of the gas concentration), the environmental parameter conditions, the features selected for evaluation, and the classifiers utilized to map the extracted features to class labels (Vergara et al., 2009a). A more comprehensive formulation addressing all of these options would almost certainly yield a better performance than the generic solutions presented in these works (Di Natale et al., 1995; Wilson and De Weerth, 1995; Martinelli et al., 2003; Vergara et al., 2007b,d). We held this issue, though, as an arguable position that we would address in future works.

Finally, when facing a general problem of sensor-array optimization, sensor selection, or sensor optimization, it would be worth considering the following steps:

1. In case there is domain knowledge (e.g., underlying physicochemical phenomena is known or a sensor response model is available), then consider building a better set of "*ad hoc*" features.
2. Are features commensurate? If a hybrid multi-sensor system is employed, which combines different odor sensing technologies it is likely that feature normalization will be of help.
3. If there is a need to prune the number of sensors or input features (e.g., for simplicity, data understanding, etc.) then

start by assessing features individually. This will help in understanding their relative influence on the system. Additionally, if their number is too large use a variable ranking method to implement a first step of filtering.

4. In case some response patterns within the optimization database are suspected to be meaningless or have the wrong class labels, outliers should be detected by employing the top ranking features/sensors obtained in step 3.
5. In case there is not an evident option for the classifier model to be used, start by trying first a linear model. Following the ranking of step 3, construct a sequence of classifiers of similar nature using increasing (or decreasing) subsets of features (e.g., by implementing a forward or backward selection strategy). In case performance is matched or improved with a smaller subset, then try a non-linear classifier model with this subset. As a rule of thumb, it is better to try simple strategies first.
6. If the database available has a large number of samples and there is time enough and computational resources, then it is worth comparing several sensor/feature selection methods coupled to linear or non-linear classifier models. Consider combining filters either with wrapper or embedded approaches. Do not refrain from trying your own ideas.
7. Finally, in order to achieve a stable solution to the optimization problem and effectively improve performance, employ data resampling methods (e.g., by constructing bootstrap sets) and redo the sensor selection analysis for different sets.

ACKNOWLEDGMENTS

This work was partially funded by NATO under the Science for Peace Program grant number CBP.MD.CLG 983914. A. Vergara is funded by the U.S. Office of Naval Research (ONR), contract number N00014-07-1-0741 and by the US Army Medical Research and Materiel Command and the United States Army Research Institute of Environmental Medicine (USARIEM), contract number W81XWH-10-C-0040 in collaboration with Elintrix Inc. E. Llobet is supported by the Spanish Ministry of Science and Innovation and the Catalan Agency for Research under the grant numbers TEC 2009-07107 and 2009-SGR-789, respectively. The authors also acknowledge Joanna Zytkowicz for reading and revising the manuscript.

REFERENCES

Alkasab, T. K., White, J., and Kauer, J. S. (2002). A computational system for simulating and analyzing arrays of biological and artificial chemical sensors. *Chem. Senses* 27, 261–275.

Bhandawat, V., Olsen, S. R, Gouwens, N. W., Schlief, M. L., and Wilson, R. I. (2007). Sensory processing in the *Drosphila* antennal lobe increases reliability and separability of ensemble odor representations. *Nat. Neurosci.* 10, 1474–1482.

Barsan, N., Schweizer-Berberich, M., and Göpel, W. (1999). Fundamental and practical aspects in the design of nanoscaled SnO2 gas sensors: a status report. *Fresenius J. Anal. Chem.* 365, 287–304.

Cavicchi, R. E., Suehle, J. S., Kreider, K. G., Gaitan, M., and Chaparala, P. (1996). Optimized temperature-pulse sequences for the enhancement of chemically specific response patterns from micro-hotplates gas sensors. *Sens. Actuators B Chem.* 33, 142–146.

Choi, N. H., Shim, C. H., Song, K. D., Lee, D. S., Huh, J. S., and Lee, D. D. (2002). Classification of workplace gases using temperature modulation and two SnO2 sensing films on substrate. *Sens. Actuators B Chem.* 86, 251–258.

Clifford, P. K., and Tuma, D. T. (1983). Characteristics of semiconductor gas sensor II, Transient response to temperature change. *Sens. Actuators B Chem.* 3, 233–254.

Comini, E., Baratto, C., Faglia, G., Ferroni, M., Vomiero, A., and Sberveglieri, G. (2009). Quasi-one dimensional metal oxide semiconductors: preparation, characterization and application as chemical sensors. *Prog. Mat. Sci.* 54, 1–67.

Corcoran, P., Lowery, P., and Anglesea, J. (1998). Optimal configuration of a thermally cycled gas sensor array with neural network pattern recognition. *Sens. Actuators B Chem.* 48, 448–455.

Davis, L. (1991). *The Handbook of Genetic Algorithms.* New York, NY: Van Nostrand Reinhold.

Denzler, J., and Brown, C. M. (2002). Information theoretic sensor data selection for active object recognition and state estimation. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 145–157.

Di Natale, C., Marco, S., Davide, F., and D'Amico, A. (1995). Sensor array calibration time reduction by dynamic modeling. *Sens. Actuators B Chem.* 24, 578–583.

Dickinson, T. A., White, J., Kauer, J. S., and Walt, D. R. (1996). A chemical-detecting system based on a cross-reactive optical sensor array. *Nature* 382, 697–700.

Eklöv, T., Mårtensson, P., and Lundström, I. (1997). Enhanced selectivity of MOSFET gas sensors by systematical analysis of transient parameters. *Anal. Chim. Acta* 353, 291–300.

Floreano, D., Kato, T., Marocco, D., and Sauser, E. (2004). Coevolution of active vision and feature selection. *Biol. Cybern.* 90, 218–228.

Fort, A., Gregorkiewitz, M., Machetti, M., Rocchi, S., Serrano, B., Tondi, L., Ulivieri, N., Vignoli, V., Faglia, G., and Comini, E. (2002). Selectivity enhancement of SnO2 sensors by means of operating temperature modulation. *Thin Solid Films* 418, 2–8.

Fort, A., Machetti, M., Rocchi, S., Serrano, B., Tondi, L., Ulivieri, N., Vignoli, V., and Sberveglieri, G. (2003). Tin oxide gas sensing: comparison among different measurement techniques for gas mixture classification. *IEEE Trans. Instrum. Meas.* 52, 921–926.

Franke, M. E., Koplin, T. J., and Simon, U. (2006). Metal and metal oxide nanoparticles in chemiresistors: does the nanoscale matter? *Small* 2, 36–50.

Freund, M. S., and Lewis, N. S. (1995). A chemically diverse, conducting polymer-based "electronic nose". *Proc. Natl. Acad. Sci. U.S.A.* 92, 2652–2656.

Gardner, J. W., and Bartlett, P. N. (1996). Performance definition and standardization of electronic noses. *Sens. Actuators B Chem.* 33, 60–67.

Gardner, J. W., and Bartlett, P. N. (1999). *In Electronic Noses: Principles and Applications.* Oxford: Oxford University Press.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin.

Godfrey, K. (1993). *Perturbation Signals for System Identification.* Prentice Hall International, 39–49, 181–187, 321–347.

Göpel, W. (1985). Chemisorption and charge transfer at ionic semiconductor surfaces: VCH implications in designing gas sensors. *Prog. Surf. Sci.* 20, 1, 9–103.

Göpel, W. (1988). Chemical imaging I: concepts and visions for electronic and bioelectronic noses. *Sens. Actuators B Chem.* 52, 125–142.

Göpel, W., Hesse, J., and Zemel, J. N. (eds). (1991). *Sensors, A Comprehensive Survey*, Vol. 2. Weinheim: VCH.

Gosangi, R., and Gutierrez-Osuna, R. (2009). "Active chemical sensing with partially observable Markov decision processes," in *Proceedings of International Symposium on Olfaction and Electronic Noses (ISOEN 2009)*, Brescia.

Gosangi, R., and Gutierrez-Osuna, R. (2010). Active temperature programming for metal-oxide chemoresistors. *IEEE Sens. J.* 10, 1075–1082.

Gurlo, A. (2010). Nanosensors: does crystal shape matter? *Small* 6, 2077–2079.

Gutierrez-Osuna, R., Nagle, H. T., and Schiffman, S. S. (1999). Transient response analysis of an electronic nose using multi-exponential models. *Sens. Actuators B Chem.* 61, 170–182.

Guyon, I., and Elisseeff, A. (2003). An introduction to variable and feature selection. *J. Mach. Learn. Res.* 3, 1157–1182.

Huang, X., Liu, J., Shao, D., Pi, Z., and Yu, Z. (2003). Rectangular mode of operation for detecting pesticide residue by using a single SnO2 based gas sensor. *Sens. Actuators B Chem.* 96, 630–635.

Kullback, S., and Leibler, R. A. (1951). On information and sufficiency. *Ann. Math. Stat.* 22, 1, 79–86.

Kunt, T. A., McAvoy, T. J., Cavicchi, R. E., and Semancik, S. (1998). Optimization of temperature programmed sensing for gas identification using micro-hotplate sensors. *Sens. Actuators B Chem.* 53, 24–43.

Martinelli, E., Falconi, C., D'Amico, A., and Di Natale, C. (2003). Feature extraction of chemical sensors in phase space. *Sens. Actuators B Chem.* 95, 132–139.

Maziarz, W., and Pisarkiewicz, T. (2008). Gas sensors in a dynamic operation mode. *Meas. Sci. Technol.* 19, 055205.

Moseley, P. T., and Tofield, B. C. (eds). (1987). *Solid-state Gas Sensors. The Adam Hilger Series on Sensors*, Bristol.

Muezzinoglu, M., Vergara, A., Huerta, R., and Rabinovich, M. (2010). A sensor conditioning principle for odor identification. *Sens. Actuators B Chem.* 146, 472–476.

Muezzinoglu, M., Vergara, A., Huerta, R., Rulkov, N., Rabinovich, M., Selverston, A., and Abarbanel, H. (2009) Acceleration of chemosensory information processing using transient features. *Sens. Actuators B Chem.* 137, 507–512.

Nakamoto, T., Okazaki, N., and Matsushita, H. (1995). Improvement of optimisation algorithm in active gas/odour sensing system. *Sens. Actuators A Phys.* 50, 191–196.

Nakamoto, T., Okazaki, N., and Morizumi, T. (1997). High speed active gas/odor sensing system using adaptive control theory. *Sens. Actuators B Chem.*, 41, 183–188.

Nakata, S., Akakabe, S., Nakasuji, M., and Yoshikawa, K. (1996). Gas sensing based on a nonlinear response: discrimination between hydrocarbons and quantification of individual components in a gas mixture. *Anal. Chem.* 68, 2067–2072.

Nakata, S., Kaneda, Y., Nakamura, H., and Yoshikawa, K. (1991). Detection and quantification of CO gas based on the dynamic response of a ceramic sensor. *Chem. Lett.* 20, 1505–1508.

Nakata, S., Kaneda, Y., and Yoshikawa, K. (1992). Novel strategy to develop chemical sensors based on nonlinear dynamics-intelligent gas sensor. *Sens. Mater.* 4, 101–110.

Nakata, S., Ozaki, E., and Ojima, N. (1998). Gas sensing based on the dynamic nonlinear responses of a semiconductor gas sensor: dependence on the range and frequency of a cyclic temperature change. *Anal. Chim. Acta* 361, 93–100.

Niebling, G., and Müller, R. (1995). Design of sensor arrays by use of an inverse feature space. *Sens. Actuators B Chem.* 25, 781–784.

Paletta, L., and Pinz, A. (2000). Active object recognition by view integration and reinforcement learning. *Robot. Auton. Syst.* 31, 71–86.

Papadimitriou, C. H., and Tsitsiklis, J. N. (1987). The complexity of Markov decision processes. *Math. Oper. Res.* 12, 441–450.

Pearce, T. C., and Sanchez-Montañes, M. (2003). "Chemical sensor array optimisation: geometric and information theoretic approaches," in *Handbook of Artificial Olfaction Machines*, Chapt. 14, eds T. C. Pearce, S. Schiffman, H. T. Nagle, and J. Gardner (Weinheim: Wiley-VCH), 347–376.

Persaud, K., and Dodd, G. (1982). Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature* 299, 352–355.

Persaud, K. C., and Travers, P. J. (1997). *Handbook of Biosensors and Electronic Noses.* Boca Raton: CRC Press, 563–592.

Raman, B., Meier, D. C., Evju, J. K., and Semancik, S. (2009). Designing and optimizing microsensor arrays for recognizing chemical hazards in complex environments. *Sens. Actuators B Chem.* 137, 617–629.

Rodriguez-Lujan, I., Huerta, R., Elkan, C., and Santa Cruz, C. (2010). Quadratic programming feature selection. *J. Mach. Learn. Res.* 11, 1491–1516.

Ruhland, B., Becker, T., and Müller, G. (1998). Gas-kinetic interactions of nitrous oxides with SnO2 surfaces. *Sens. Actuators B Chem.* 50, 85–94.

Sears, W. M., Colbow, K., and Consadori, F. (1989a). General characteristics of thermally cycled tin oxide gas sensors. *Semicond. Sci. Technol.* 4, 351–359.

Sears, W. M., Colbow, K., and Consadori, F. (1989b). Algorithms to improve the selectivity of thermally cycled tin oxide gas sensors. *Sens. Actuators B Chem.* 19, 333–349.

Sears, W. M., Colbow, K., Slamka, R., and Consadori, F. (1990). Selective thermally cycled gas sensing using fast Fourier-transform techniques. *Sens. Actuators B Chem.* 2, 283–289.

Semancik, S., and Cavicchi, R. E. (1999). Kinetically controlled chemical sensing using micromachined structures. *Acc. Chem. Res.* 31, 5, 279–287.

Stoycheva, T., Vallejos, S., Blackman, C., Moniz, S. J. A., Calderer, J., and Correig, X. (2011). Important considerations for effective gas sensors based on metal oxide nanoneedles films. *Sens. Actuators B Chem.* doi:10.1016/j.snb.2011.10.052

Vergara, A., and Llobet, E. (2011). Feature selection versus feature compression in the building of calibration models from FTIR-spectrophotometry datasets. *Talanta.* doi:10.1016/j.talanta.2011.10.014

Vergara, A., Llobet, E., Brezmes, J., Ivanov, P., Cané, C., Gràcia, I., Vilanova, X., and Correig, X. (2007a). Quantitative gas mixture analysis using temperature-modulated micro-hotplate gas sensors: selection and validation of the optimal modulating frequencies. *Sens. Actuators B Chem.* 123, 1002–1016.

Vergara, A., Llobet, E., Martinelli, E., Di Natale, C., D'Amico, A., and Correig, X. (2007b). Feature extraction of metal oxide gas sensors using dynamic moments. *Sens. Actuators B Chem.* 122, 219–226.

Vergara, A., Llobet, E., Ramírez, J. L., Ivanov, P., Fonseca, L., Zampolli, S., Scorzoni, A., Becker, T., Marco, S., and Wöllenstein, J. (2007c). An RFID reader with onboard sensing capability for monitoring fruit quality. *Sens. Actuators B Chem.* 127, 143–149.

Vergara, A., Martinelli, E., Llobet, E., Giannini, F., D'Amico, A., and Di Natale, C. (2007d). An alternative global feature extraction of temperature modulated micro-hotplate gas sensors array using an energy vector approach. *Sens. Actuators B Chem.* 124, 352–359.

Vergara, A., Llobet, E., Brezmes, J., Vilanova, X., Ivanov, P., Gràcia, I., Cané, C., and Correig, X. (2005a). Optimized temperature modulation of micro-hotplate gas sensors through pseudo random binary sequences. *IEEE Sens. J.* 5, 1369–1378.

Vergara, A., Llobet, E., Brezmes, J., Ivanov, P., Vilanova, X., Gràcia, I., Cané, C., and Correig, X. (2005b). Optimised temperature modulation of metal oxide micro-hotplate gas sensors through multilevel pseudo random sequences. *Sens. Actuators B Chem.* 111–112, 271–280.

Vergara, A., Martinelli, E., Huerta, R., D'Amico, A., and Di Natale, C. (2011). Orthogonal decomposition of chemo-sensory cues. *Sens. Actuators B Chem.* 159, 126–134.

Vergara, A., Martinelli, E., Llobet, E., D'Amico, A., and Di Natale, C. (2009a). Optimized feature extraction for temperature-modualted gas sensors. *J. Sens.* 2009, 716316.

Vergara, A., Muezzinoglu, M. K., Rulkov, N., and Huerta, R. (2009b). "Kullback-Leibler distance optimisation for artificial chemo-sensors," in *IEEE Sensors Conference*, Christchurch.

Vergara, A., Muezzinoglu, M. K., Rulkov, N., and Huerta, R. (2010). Information theory of chemical sensors. *Sens. Actuators B Chem.* 148, 298–306.

Vergara, A., Ramirez, J. L., and Llobet, E. (2008). Reducing power consumption via a discontinuous operation of temperature-modulated micro-hotplate gas sensors: application to the logistics chain of fruit. *Sens. Actuators B Chem.* 129, 311–318.

White, J., Kauer, J. S., Dickinson, T. A., and Walt, D. R. (1996). Rapid analyte recognition in a device based on optical sensors and the olfactory system. *Anal. Chem.* 68, 2191–2202.

Wilson, D. M., and De Weerth, S. P. (1995). Odor discrimination using steady-state and transient characteristics of tin-oxide sensors. *Sens. Actuators B Chem.* 28, 123–128.

Wlodek, S., Colbow, K., and Consadori, F. (1991). Kinetic model of thermally cycled tin oxide gas sensor. *Sens. Actuators B Chem.* 3, 123–127.

Yamazoe, N. (2005). Toward innovations of gas sensor technology. *Sens. Actuators B Chem.* 108, 2–14.

Zaromb, S., and Stetter, J. R. (1984). Theoretical basis for identification and measurement of air contaminants using an array of sensors having partly overlapping selectivities. *Sens. Actuators B Chem.* 6, 225–243.

Zierler, N. (1959). Linear recurring sequences. *J. Soc. Ind. Appl. Math.* 7, 31–48.

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.

Received: 05 October 2011; paper pending published: 20 October 2011; accepted: 08 December 2011; published online: 04 January 2012.

Citation: Vergara A and Llobet E (2012) Sensor selection and chemo-sensory optimization: toward an adaptable chemo-sensory system. Front. Neuroeng. ***4****:19. doi: 10.3389/fneng.2011.00019*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 12 January 2012
doi: 10.3389/fneng.2011.00020

Mobile robots for localizing gas emission sources on landfill sites: is bio-inspiration the way to go?

Victor Hernandez Bennetts[1], Achim J. Lilienthal[1], Patrick P. Neumann[2] and Marco Trincavelli[1]*

[1] *Center for Applied Autonomous Sensor Systems, School of Science and Technology, Örebro University, Örebro, Sweden*
[2] *BAM Federal Institute for Materials Research and Testing, Berlin, Germany*

Edited by:
Ramon Huerta, University of California San Diego, USA

Reviewed by:
Sandro Carrara, Ecole Polytechnique Fédérale de Lausanne, Switzerland
Francisco B. Rodriguez, Universidad Autonoma de Madrid, Spain

****Correspondence:***
Marco Trincavelli, Center for Applied Autonomous Sensor Systems, School of Science and Technology, Örebro University, Örebro SE-70182, Sweden.
e-mail: marco.trincavelli@oru.se

Roboticists often take inspiration from animals for designing sensors, actuators, or algorithms that control the behavior of robots. Bio-inspiration is motivated with the uncanny ability of animals to solve complex tasks like recognizing and manipulating objects, walking on uneven terrains, or navigating to the source of an odor plume. In particular the task of tracking an odor plume up to its source has nearly exclusively been addressed using biologically inspired algorithms and robots have been developed, for example, to mimic the behavior of moths, dung beetles, or lobsters. In this paper we argue that biomimetic approaches to gas source localization are of limited use, primarily because animals differ fundamentally in their sensing and actuation capabilities from state-of-the-art gas-sensitive mobile robots. To support our claim, we compare actuation and chemical sensing available to mobile robots to the corresponding capabilities of moths. We further characterize airflow and chemosensor measurements obtained with three different robot platforms (two wheeled robots and one flying micro-drone) in four prototypical environments and show that the assumption of a constant and unidirectional airflow, which is the basis of many gas source localization approaches, is usually far from being valid. This analysis should help to identify how underlying principles, which govern the gas source tracking behavior of animals, can be usefully "translated" into gas source localization approaches that fully take into account the capabilities of mobile robots. We also describe the requirements for a reference application, monitoring of gas emissions at landfill sites with mobile robots, and discuss an engineered gas source localization approach based on statistics as an alternative to biologically inspired algorithms.

Keywords: mobile robotics, mobile robot olfaction, landfill surveillance, biologically inspired robots

1. INTRODUCTION

Gas-sensitive mobile robots are valuable instruments for addressing tasks like detection of gas leaks, search for explosives, surveillance, or exploration of areas where hazardous gases might be present. A concrete example of these application fields that is gaining particular interest among the EU authorities is monitoring of gaseous emissions in landfill sites (Scharff, 2008). Landfill emissions account for 2% of the total greenhouse gases (GHG) released by human activity (Bogner et al., 2007). The GHG released from landfill sites are mainly methane (CH_4) and carbon dioxide (CO_2), and to a minor extent mixtures of oxygen (O_2), nitrogen (N_2), and hydrogen (H_2). Also, poisonous gases like hydrogen sulfide (H_2S) can be released in landfills. In particular, methane produced from solid waste is a biogas that can be used as an alternative energy source (Scharff, 2008). According to Atleverket, the agency responsible of waste management in the province of Örebro, Sweden, 18000 MWh are produced yearly from the biogas that is obtained from the landfill sites in Örebro. It is important to notice that methane leaks occur frequently even in landfills that have been closed for decades. These leaks are difficult to detect and with current monitoring technologies it can take weeks before a leak is detected and localized. Ultimately, this turns into a waste of resources and in a substantial emission of GHG. The delay and inaccuracy in the leak detection are mainly due to the sparsity, both spatial and temporal, of collected samples. Currently, a landfill operator is required by law to collect one sample a month for a few predetermined locations. Mobile robotics can make a significant contribution in this area by providing versatile systems for autonomous monitoring of diverse environments. Robotic solutions can adaptively collect sensor measurements, cooperate with other systems, and provide useful indications to landfill operators. Compared to human operators, mobile robots have the advantage to carry out the required repetitive measurement procedure without suffering from fatigue and therefore, they can perform measurements with a much denser spatio-temporal granularity. Moreover, the use of an automated monitoring platform can minimize the exposure of human operators to hazardous compounds like, for example, H_2S. In addition, the mobile robots that carry the sensors offer the required accurate localization and computational resources to compute for example, on-line gas distribution models. This enables the possibility to decide which locations to observe next based on the current model (Neumann et al., 2012).

A landfill monitoring robot should be able to perform two major tasks. (1) Serve as an autonomous and flexible system that

can explore large areas in order to measure biogas concentrations (e.g., CO_2 and CH_4). (2) From the acquired measurements, it should be able to provide the landfill operators with useful information such as gas distribution maps and locations of possible gas leaks. Strictly, for none of the two tasks there is a direct biological example to mimic. However, these two complex tasks can be broken down into specific subtasks like discriminating among the different gases, estimating the location of multiple gas sources and building gas distribution maps. For what concerns odor discrimination, many works can be found in literature that formulate mathematical models of the olfactory pathways that process the signal coming from the olfactory receptors (Cleland and Linster, 2005). While, to the knowledge of the authors, the construction of a gas distribution map has no direct biological counterpart, insects have an outstanding ability to find distant sources of odors (Cardé and Willis, 2008). Insects locate odor sources by tracking wind-borne odor plumes to their emission source. The majority of robots that have been proposed for gas source localization take inspiration from insects and try to track an odor plume up to its source, where they declare the end of their task (Lilienthal and Duckett, 2003; Lilienthal et al., 2003, 2006b; Kowadlo and Russell, 2008). The animals that have inspired most of the robotics research for odor plume tracking are:

Moths, which use odor localization to find mates (Kuwana et al., 1999; Ishida et al., 2001; Pyk et al., 2006).
Lobsters, which use odor localization to locate food (Grasso et al., 1998).
Escherichia Coli, which use odor localization to locate nutrients (Russell et al., 2003).
Dung Beetles, which use odor localization to find hatching niches, habitation, and food (Russell et al., 2003).

The bio-inspired gas source tracking algorithms that have been implemented on mobile robots are based on two principles, namely *chemotaxis* and *anemotaxis*. *Chemotaxis* refers to a mechanism in which the movement of an organism (or robot) is determined by the distribution of chemical compounds, most often by the concentration gradient of one or more volatiles. *Anemotaxis* instead refers to a mechanism in which the movement of an organism (or robot) is determined by the perceived airflow (air can be generalized to fluid).

More recently Vergassola et al. (2007) proposed *infotaxis*, which is a search strategy based on probability and information theory. It was designed for addressing the gas source localization problem in uncontrolled indoor or outdoor environments characterized by a high Reynolds number and thus dominated by turbulence (Roberts and Webster, 2002). Instead of using concentration and flow gradients, *infotaxis* models the location of the source as a probability distribution derived from previously collected measurements and the next actions of the robot (i.e., move to a neighboring location or standing still) are decided based on a minimum entropy criterion.

To the knowledge of the authors, only two works have evaluated the concepts presented by *infotaxis* with physical experiments. In Moraud and Martinez (2010), the authors assess the robustness and reliability of *infotaxis* for localizing a heat source, arguing that heat has dispersion properties very similar to chemical compounds. However, temperature sensors have a much faster dynamics than gas sensors. In addition, an artificial airflow was induced with a fan, which simplifies the gas source localization problem even further. In Lochmatter (2010), the authors evaluated, among other source location algorithms, a statistical approach that follows the same principles as *infotaxis*. The authors performed their experiments inside a 18 m × 4 m wind tunnel with an ethanol gas source and a robot equipped with a commercial gas sensor. The experiments were carried out under laminar flow conditions.

In general, most of the work in mobile robot olfaction has been developed under simplified assumptions. The most common assumptions on which works in literature are based are steady constant airflow and the presence of a single gas source emitting a known chemical compound at a constant rate. Moreover, in most of the cases the exploration area is of limited size, the airflow is artificially modified and the robot's starting position is located downwind with respect to the gas source, making it easy for the robot to collect odor cues. Obviously, those assumptions do not hold in a scenario like a landfill. First of all, a landfill is an area of considerable size, where finding gas traces that can initiate a plume tracking algorithm is difficult. Therefore, initial exploration strategies for gas finding cannot be trivially formulated. Moreover, more than one gas leak emitting different compounds (at variable rates) may be present at the same time. Finally, the airflow in an open outdoor environment like a landfill is dominated by turbulent advection (Shraiman and Siggia, 2000). Turbulent airflow disperses the gas plume creating a complex structure of gas patches with different concentration levels. Furthermore, advection can create areas of high concentration away from the location where the gas was released.

The main contribution of this paper is to highlight the weaknesses of state-of-the-art bio-inspired algorithms for gas source localization that aim to directly reproduce insect behavior. The argumentation brought at support is twofold: first, the sensing mechanisms available to robotic systems are completely different from biological receptors, and second, the chaotic environmental properties of natural environments do not allow the formation of a steady odor plume that would lead a robot that implements a form of bio-inspired anemotaxis to the gas source. Moreover, for localizing a gas source a mobile robot does not necessarily need to travel toward it tracking the odor plume. Indeed, the robot can collect measurements in locations far away from the gas source and still be able to infer the position of the emission source.

We provide observations collected in four different experimental areas with three different platforms in order to support our claims. We then use an engineered approach as an alternative to biologically inspired gas localization algorithms. To conclude, we argue that isolated principles of animal behavior can help us to understand the task of gas source localization better, provided that we carefully take into account the limited mobility of a robot and the difference between biological receptors and the sensors that are available to a robot.

2. MATERIALS AND METHODS

In order to investigate the characteristics of measurements collected by mobile robots in natural environments, we perform experiments in four different locations with three different robotic

platforms. The different locations have been chosen in order to have a wide spectrum of possible environmental conditions, ranging from a closed isolated room, to a long corridor with open ends and two different courtyards. Also, the selection of the robotic platform aims at investigating different sensor configurations imposed by the constraints of the different platforms. In particular, the experiments were performed with two wheeled robots and one flying quadrocopter. Wheeled robots have a high payload, a long battery life, and they can have large computational resources onboard. On the other hand flying platforms have less restrictions regarding mobility compared to wheeled robots but typically they have limited payload and battery life. Another shortcoming of helicopter platforms is that the action of their rotors can influence significantly the airflow, modifying the original gas distribution. In all the experiments the robots were following a predefined trajectory that covered the area of inspection. The robots were stopping at regular intervals in order to collect series of measurements. No attempt was made to move the robot toward the location of the gas source but instead, data have been collected over all the area of inspection and then, gas distribution maps have been produced at the end of the experiments. As we will see in Section 3, gas distribution maps can provide good indications on the location of the gas source. In the next subsections we will first introduce the robotics platforms (Section 2.1), then describe the experimental environments (Section 2.2), and finally illustrate the sensing modalities employed for chemical (Section 2.3) and airflow (Section 2.4) sensing.

2.1. ROBOTIC PLATFORMS

2.1.1. Outdoor wheeled robot

The first robot, shown in **Figure 1A** is an ATRV-JR all terrain robot equipped with a laser range scanner (SICK LMS 200) used for localization, an electronic nose, and an anemometer. The electronic nose comprises six gas sensors (five metal oxide and one electrochemical) enclosed in an aluminum tube. This tube is horizontally mounted at the front side of the robot at a height of 0.1 m over the ground. The electronic nose is actively ventilated through a fan that creates a constant airflow toward the gas sensors. Thus, it lowers the effect of external airflow and the movement of the robot on the sensor response and guarantees a continuous exchange of gas in situations with very low external airflow. The gas sensors used in the electronic nose are listed in **Table 1**. The ultrasonic anemometer used to measure the airflow is a Young 81000 with a range from 0.02 up to 40 m/s and a resolution of 0.01 m/s. The placement of the anemometer had to be a compromise between the desire to measure the airflow as close to the gas sensors and as undisturbed as possible. It was finally placed above the top of the robot in order to minimize the influence of the fan of the electronic nose and the body of the robot itself. The robot software is based on the Player robot server (Gerkey et al., 2003), a control interface that simplifies access to standard robot sensors and actuators and provides implementations of standard algorithms. In particular, adaptive Monte Carlo localization (amcl driver), VFH+ obstacle avoidance (vfh driver), and the wavefront path planner (wavefront driver) were used for localization, local, and global path planning, respectively. The localization module implemented in the amcl driver uses the odometry and the laser scanner readings in order to localize the robot on a map provided to the algorithm at the startup.

2.1.2. Indoor wheeled robot

The second robot used in our experiments is a Pioneer P3-DX (MobileRobots), shown in **Figure 1B**. The Pioneer P3-DX is a smaller platform compared to the ATRV-JR. It has been selected due to its better maneuverability, a desirable characteristic while performing experiments in a small room. As in the case of the ATRV-JR, the robot is equipped with a laser scanner and it runs the Player robot server for performing localization and navigation tasks. For what concerns gas sensing devices, the robot was equipped with three MOX gas sensors (TGS2620, Figaro Engineering) positioned at different heights (0.571, 0.393, and 0.199 m on the ground). Moreover, the robot was equipped with two e-noses containing the 11 sensors listed in **Table 1** and a photo ionization (PID) gas detector (ppbRAE 3000, RAE Systems). The e-noses and the PID inlet were mounted 0.065 m over the ground. The photo ionization gas detector shows quick response to a wide range of gases, and provides calibrated readings of the gas concentration (given that the chemical compound is known). The responses

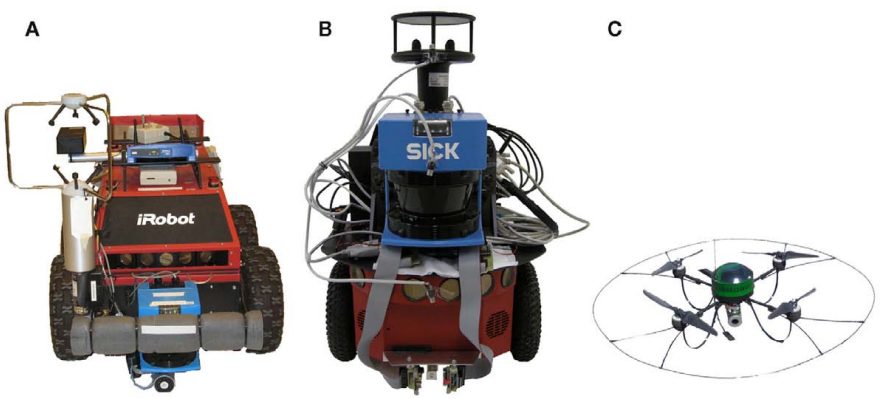

FIGURE 1 | Robots used in the presented experiments. Notice that the three pictures are not in scale. **(A)** Outdoor wheeled robot **(B)** indoor wheeled robot **(C)** outdoor flying robot.

of the metal-oxide gas sensors are slower and the sensors are not calibrated, but enable selective gas detection. The robot is also equipped with a two-dimensional ultrasonic anemometer (WindSonic, Gill). The device can measure the airflow velocity from 0 to 60 m/s with 0.01 m/s resolution. The specifications of the Gill WindSonic anemometer are very similar to the ones of the Young 81000, and the only major difference is that the WindSonic provides a 2-D reading instead of 3-D, and is much smaller.

2.1.3. Outdoor flying robot

The third robot that we used for our tests is the AR100-B micro-drone developed by AirRobot GmbH & Co. The micro-drone was modified by the Federal Institute for Materials Research and Testing (BAM, Germany) to incorporate gas-sensitive devices as payload. The sensors mounted on the micro-drone are listed in **Table 1**. The AR100-B (**Figure 1C**) is a highly maneuverable and compact platform. With a diameter of 1 m and a weight of approx. One kilogram, it supports up to 200 g of payload and its *LiPo* battery can provide a maximum flight time of about 20–30 min. The flight control relies on an on-board Inertial Measurement Unit (IMU) that comprises a three axis accelerometer and a three axis rotation rate sensor. The IMU is also used along with a GPS unit and a compass for localization purposes. Communication with the ground station is established through a 2.4-GHz RF link in which the data packets sent can include steering instructions or data coming from the payload and micro-drone sensors. Due to the restrictions imposed by the platform, the micro-drone doesn't carry an anemometer. Instead, wind measurements are estimated by fusing the different on-board sensing modalities using the wind triangle approach explained later in Section 2.4. By using this wind sensing approach, it is possible to measure windflows in the range of 0–8 m/s.

2.2. EXPERIMENTAL ENVIRONMENTS

A first set of three experiments was conducted in a 5-m × 5-m × 5-m × 2-m closed room shown in **Figure 2A**. Although no artificial airflow was induced, a weak circulating airflow field (0.01–0.03 m/s) was formed in the room by natural convection. Ethanol and 2-propanol vapors were used as detection targets, and were released from two tubes at a constant flow rate (0.2 l/min). In this set of experiments, the robot was programmed to move along a predefined spiral path that covered the whole experimental area. The robot was stopping at regular intervals for data collection. The reason for stopping the robot at each waypoint to collect wind measurements is due to the difficulty in compensating for the movement of the robot on the anemometer readings. At each measurement point, the sensor data were recorded for 30 s and the sensors were sampled at 4 Hz. A total of three experimental trials were conducted in this scenario.

Once the experiments in the closed room have been completed, experiments in less controlled environments have been carried out. The second location chosen was a section of a long corridor with open ends and a high ceiling, shown in **Figure 2B**. The area covered by the trajectory of the robot was approximately 14 m × 2.0 m. There was more disturbance in this scenario caused by people passing by and the opening of doors and windows during the run of the experiment. The gas source was a cup full of ethanol placed on the floor in the middle of the investigated corridor segment. A total of five trials were carried out in this experimental configuration. Moreover, the outdoor scenario shown in **Figure 2C** was considered. Here, four experimental trials were carried out in an 8-m × 8-m region that is part of a much bigger open area. Again, the gas source was a cup of ethanol placed in the middle of this area. In these two experimental locations, the robot followed a predefined sweeping trajectory

Table 1 | Gas sensors used in the electronic noses mounted on the three robotic platforms.

Gas sensor model	ATRV-JR	Pioneer P3-DX	AR100-B
Figaro TGS 2600	2	2	1
Figaro TGS 2602	1	1	1
Figaro TGS 2611	1	1	1
Figaro TGS 2620	1	1	1
Figaro TGS 4161	1	–	1
e2v Mics 2610	–	1	–
e2v Mics 2710	–	1	–
e2v Mics 5521	–	2	–
e2v Mics 5121	–	1	–
e2v Mics 5135	–	1	–

Please notice that the sensors listed are all MOX sensors except the TGS 4161 which is an electrochemical sensor used for CO_2 detection.

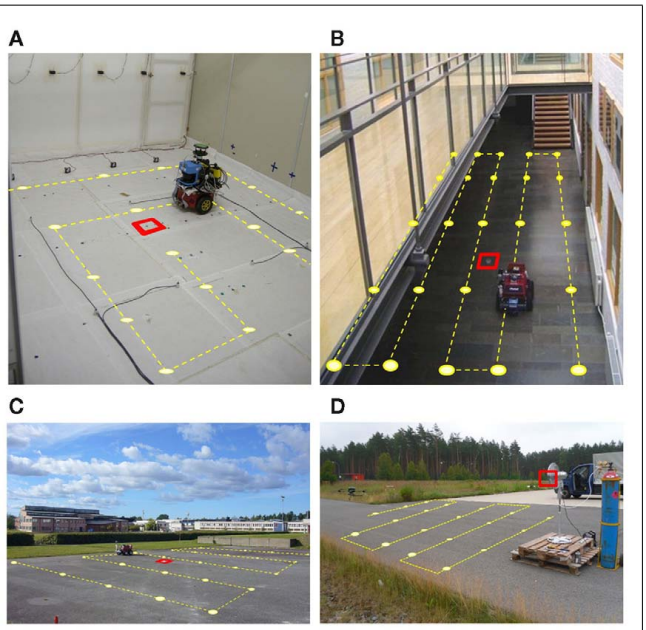

FIGURE 2 | Experimental locations considered in this work. The dashed line displays the path followed by the robot and the dots indicate the points where the robot was stopping for collecting measurements. The squares denote the actual location of the gas sources. **(A)** Closed room **(B)** Orebro University corridor **(C)** Orebro University courtyard **(D)** BAM courtyard.

covering the area of interest, using a fixed starting point. Along its path, the robot stopped at predefined positions and carried out a sequence of measurements on the spot for 10 s (outdoors) and 30 s (for both indoor locations). The predefined sweeping motion was performed once in each directions and the robot was driven at a maximum speed of 0.05 m/s in between the stops. In these two scenarios the measurements were recorded at a frequency of 1 Hz.

A set of five additional trials were conducted in the outdoor environment shown in **Figure 2D**, with an electronic nose mounted on the micro-drone previously described in Section 2.1. A CH_4 (99.5%-pure methane) gas cylinder was placed in a 14-m × 14-m area and at each trial, the valve of the cylinder was let open to release CH_4 in the environment. In order to spread the analyte away from the cylinder, an AC fan was placed near the odor outlets. The air current introduced by the fan also prevented the CH_4 to immediately rise up to the atmosphere when released. The micro-drone was programmed to explore the experimental area following a sweeping trajectory, starting from a remote location and moving at 1 m/s toward the gas cylinder. Data samples were acquired at a sampling frequency of 8 Hz and transmitted down to the ground stating using the micro-drone's RF link. At each measuring position the micro-drone stopped to take gas concentration measurements for about 20 s.

2.3. CHEMO SENSING

The metal-oxide (MOX) gas sensors are by far the most widely used in electronic nose applications as well as in mobile robotics olfaction. The most prominent reasons for this are that they are commercially available, have a relatively fast response and a higher sensitivity than most other sensing technologies. MOX gas sensors are conductometric sensors, that means that a change in the conductance of the oxide is measured when a gas interacts with the sensing surface. The logarithm of the change in resistance over a certain range is approximately linearly proportional to the logarithm of the concentration of the gas (Ihokura and Watson, 1994). There are two types of MOX sensors: *n*-type (SnO_2, ZnO), which respond to reducing gases like H_2, CH_4, CO, C_2H_5, C_2H_5OH, $(CH_3)_2CHOH$, or H_2S and *p*-type (NiO, CoO) which respond to oxidizing gases like O_2, NO_2, and Cl_2 (Janata, 2009). The response of a MOX sensor results from chemosorption and redox reactions at the surface. Since the rate of such reactions is dependent on the temperature, it is clear that the temperature of the sensing surface considerably affects the sensor characteristics (Ihokura and Watson, 1994). Typical temperatures for the sensing surface of MOX sensors lie between 300 and 500°C. Selectivity is obtained either by doping the sensing surface with different additives or by setting different operating temperatures. It has also been demonstrated that introducing a dynamic operating temperature further enhances the selectivity of the sensor (Ihokura and Watson, 1994). In addition, Gas discrimination with MOX sensors on a mobile robot has been analyzed in Trincavelli et al. (2009) and Trincavelli (2011). **Figure 3** shows a schematic of a MOX sensor. R_H and R_s are respectively the heater and the sensor resistances, while R_L is the load resistance that is applied in series to R_s in order to be able to read it. V_H is the voltage applied to the heating resistance and it is proportional to the operating temperature,

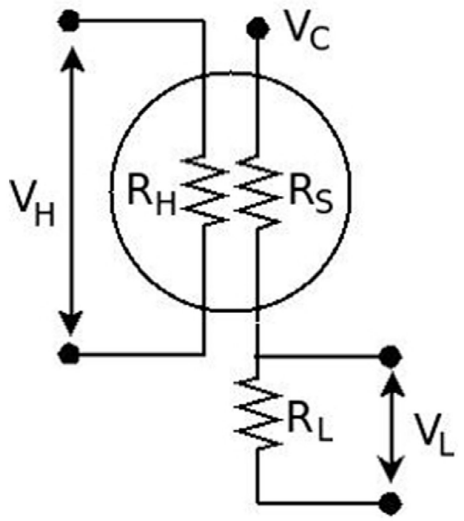

FIGURE 3 | Electrical schema of a MOX gas sensor.

V_C is the reference voltage for the measurement, while V_L is the voltage drop on R_L. In order to calculate the value of the sensor resistance (inverse of the sensor conductance – the quantity that changes when the sensor responds) the following formula is applied:

$$R_S = \frac{V_C - V_L}{V_L} \times R_L \quad (1)$$

Another gas sensor that is gaining popularity in the mobile robotics olfaction community is the photo ionization detector (PID). A PID is an ion detector which uses high-energy photons, typically in the ultraviolet range (UV), to break gas molecules into positively charged ions. As a compound enters the PID it is ionized when it absorbs high-energy UV light. In commercial PID detectors the UV light is normally provided with a 10.6-eV UV lamp. The UV light excites the molecules, which temporarily lose an electron, and thus become positively charged ions. The ions produce an electric current, which is the signal output from the detector. According to the manufacturer of the PID used in our experiments (ppbRAE 3000: Portable VOC Monitor for ppb Measurement, Available at: http://www.raesystems.com/products/ppbrae-3000), the output signal is linearly proportional to the concentration of the chemical compound being analyzed. As a standalone detector PIDs are broad band detectors and are not selective, as these may ionize everything with an ionization energy less than or equal to the lamp output. Unlike MOX gas sensors, if the chemical compound is known, PIDs provide true concentration measurements. Moreover the response dynamics of PIDs is much quicker compared to the one of MOX sensors. Two of the main drawback of PID gas sensors compared to MOX gas sensors are the high price and the considerable weight (738 g for the ppbRAE that we use in our experiments), that makes them unsuitable for platforms with a limited payload. A viable solution, that we adopt for our indoor wheeled robot is to use both sensor modalities in order to try to combine the advantages of both while limiting the shortcomings.

2.4. WIND SENSING

Wind information can be of high importance for gas-sensitive robots. For example, the steering trajectories of *anemotaxis* plume tracking algorithms are based on wind measurements. Moreover, more accurate gas distribution models can be obtained by considering the advective influence of local airflow (Reggente and Lilienthal, 2009). Anemometers based on ultrasonic measurements are a convenient solution for robotics applications due to of their high resolution, wide measurement range and their relatively compact size compared to anemometers based on mechanical principles (e.g., windmill and cup anemometers).

The basic operational principle of an ultrasonic anemometer can be explained with the schematic shown in **Figure 4A**. Pairs of piezoelectric transmitters and transducers are placed at locations s_1 and s_2 separated by a distance L. The oscillator at s_1 produces a sonic pulse that reaches the transducer at s_2 in a time of flight t_{12} while a second pulse travels from s_2 to s_1 in t_{21}. From Equation 2, the wind speed and direction can be estimated. When wind conditions are negligible, $t_{12} \approx t_{21}$ and therefore $|u| \approx 0$. If wind is blowing toward a location i, the sonic pulse emitted from a location j will travel faster than the pulse emitted from i and therefore $t_{ji} < t_{ji}$. The wind direction is then inferred from the sign of u. By placing additional orthogonal arrays of transducers/oscillators, 2-D and 3-D wind information can be acquired.

$$u = \frac{L}{2} \times \left[\frac{1}{t_{ij}} - \frac{1}{t_{ji}} \right] \tag{2}$$

While ultrasonic anemometers are a reliable means to acquire wind information, alternative approaches have been developed for platforms with particular restrictions. Specifically, the micro-drone described in Section 2.1 has a limited payload capacity which imposes weight and size constraints for on-board equipment and furthermore, the turbulence and vibrations caused by the micro-drone's rotors can disrupt the anemometer readings.

By fusing different sensing modalities, it is nevertheless possible to estimate the wind information. Neumann et al. (2010) proposed a method that can be used by aerial robots to estimate the wind vector $\vec{u}$. The authors used data coming from the embedded sensors of their micro-drone to compute the parameters of the wind triangle shown in **Figure 4B**. The ground vector $\vec{w}$ and its direction w_{dir} were directly obtained from the GPS readings while the flight vector $\vec{v}$ was calculated using a reference function computed from a set of wind tunnel measurements. Additionally, the orientation information coming from the on-board compass was used to compute the flight direction v_{dir} and finally, the wind vector $\vec{u}$ and direction u_{dir} were computed from the wind triangle by applying the law of cosines.

3. RESULTS

In this section, we analyze the data recorded with the three robotic platforms in the four experimental locations summarized in **Table 2**. We characterize the wind measurements to describe the environmental conditions that prevailed during the data collection in the four locations. In addition, we analyze the response dynamics in the frequency domain of a PID sensor and two MOX sensors that are commonly used in robotic olfaction. To conclude this section, we explore a non-biological approach to gas source localization.

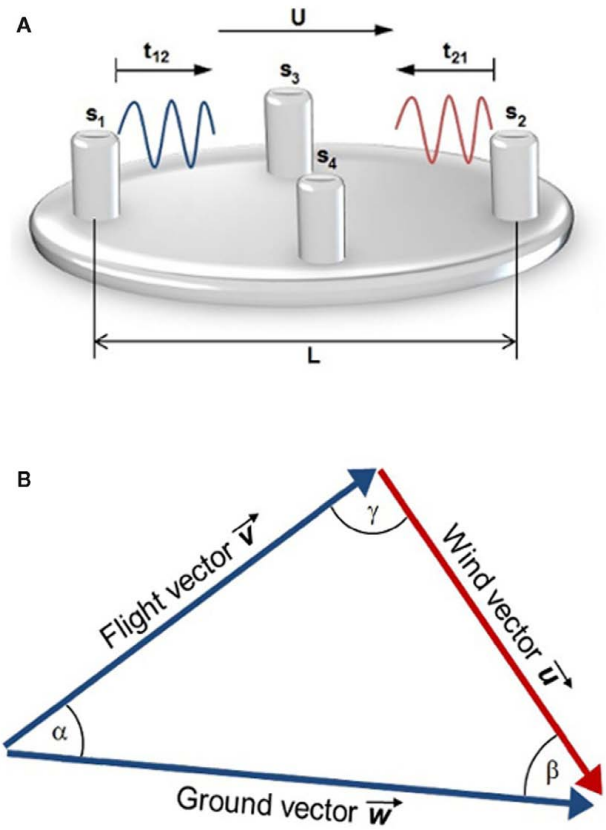

FIGURE 4 | Schematic diagrams of the wind measurement principles considered in this work. (A) Ultrasonic measurement principle **(B)** wind triangle measurement principle.

Table 2 | Measurement configurations used in the four experimental locations considered in this work.

Location	Platform	Area	Waypoint distance (m)	Gas source	Trials
Closed room	Wheeled Pioneer P3-DX	5 m × 2 m	0.5	Ethanol	3
Örebro University corridor	Wheeled ATRV-JR	14 m × 2 m	2	Ethanol	5
Örebro University courtyard	Wheeled ATRV-JR	8 m × 8 m	2	Ethanol	4
BAM courtyard	Micro-drone AR100-B	14 m × 14 m	2	Methane	5

For each of the four experimental locations, we generated a typical airflow map using the data from one of the trials. At each waypoint where the robot was stopped, a mean wind vector was computed. The size of the exploration area and the distance between consecutive waypoints are listed in **Table 2**. **Figures 5A–8A** show the exploration trajectories and the computed airflow maps. The arrow's length represents the average wind speed and the circular mean direction is represented by the arrow's orientation. Although the explored areas are not of considerable size and the measurement points are spatially dense, regularity in the wind flow direction is hardly observed. This clearly indicates that the assumption of laminar air flow does not hold in any of the four environments. Large directional fluctuations were also observed between measurements taken at a single waypoint. **Figures 5B–8B** show polar plots, which were computed from a selected waypoint (denoted by red squares in the corresponding figures) on the robot's trajectory. Each arrow in the plot is a measurement taken at the waypoint. The length is proportional to the wind speed and the arrow's direction represents the wind angle. The polar wind measurement plots show an irregular distribution of the measured wind direction, which is certainly non-Gaussian. Accordingly, the circular mean direction (denoted by a dashed red line) is not a good indicator of the wind conditions present at the waypoint. The wind distribution is further characterized by wind speed histograms that can be seen in **Figures 5C–8C**. Again, uneven speed distributions were sensed even in the indoor experimental set-ups where one might expect less distinct fluctuations.

Figure 9A displays an example of the response of the chemosensors employed in our experiments, namely a MOX gas sensors and a PID. The response of the PID is linear with respect to the chemical compound concentration and the rise and decay time constants of this sensor are symmetric and much smaller than of the MOX sensors. The PID response thus provides a good reference of the concentration the MOX gas sensor was exposed to. The plot in **Figure 9A** shows the non-linearity in the response of the MOX gas sensor, and most importantly, the slow dynamics of the MOX gas sensor. The asymmetric low-pass filtering performed by the MOX sensor is evident, especially during the long recovery of the MOX sensors. The spectra of the measurements collected with the MOX sensors and the PID are plotted in **Figure 9B**. As expected the frequency content of the signal collected with MOX sensors is much smaller than the one of a signal collected with the PID.

It has been previously reported that statistical moments can be used as indicators of gas source proximity (Lilienthal and Duckett, 2004). More specifically, the variance of a set of gas concentration measurements has been suggested as a feature that can identify the location of a source of gas (Lilienthal et al., 2006a). From the collected sensor measurements, we explored the feasibility of inferring the location of a gaseous source through the use of a gas distribution model generated by the Kernel DM + V algorithm proposed by Lilienthal et al. (2009). This algorithm is a non-parametric estimation approach that neither makes strong assumptions about the particular form of the modeled gas distribution, nor relies on expensive fluid dynamics computations to generate the model. Instead, Kernel DM + V takes a set of spatially located measurements and computes a discretized grid model where, for each cell k, a confidence value as well as distribution

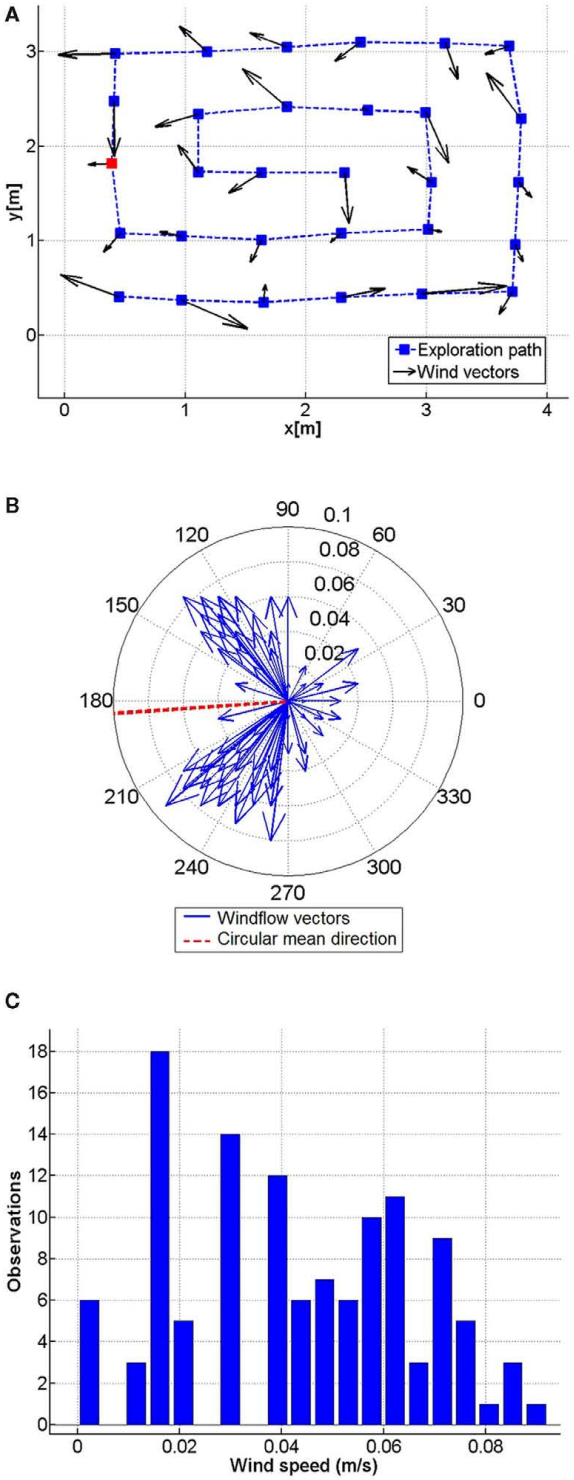

FIGURE 5 | Wind measurements collected during one experimental trial, closed room set-up. (A) Airflow map. The squares denote the points where the robot stopped to collect measurements. The dashed line denotes the robot trajectory. **(B)** Polar plot of the measurements acquired at the waypoint denoted by a red square in the airflow map. **(C)** Wind speed histogram for the measurements acquired at the waypoint denoted by a red square in the airflow map.

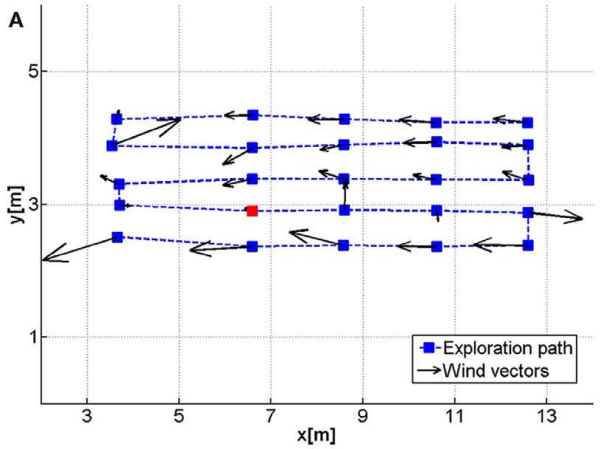

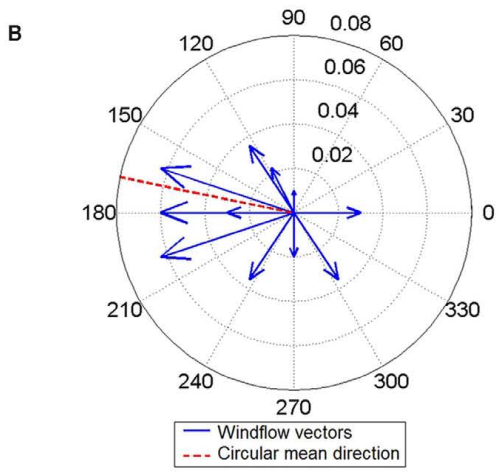

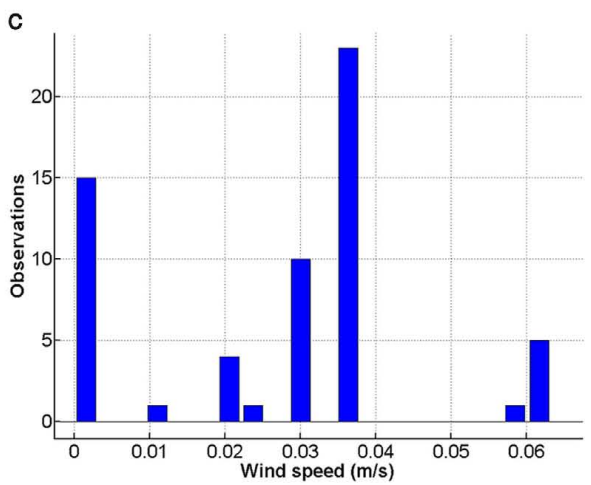

FIGURE 6 | Wind measurements collected during one experimental trial, Örebro University corridor set-up. (A) Airflow map. The squares denote the points where the robot stopped to collect measurements. The dashed line denotes the robot trajectory. **(B)** Polar plot of the measurements acquired at the waypoint denoted by a red square in the airflow map. **(C)** Wind speed histogram for the measurements acquired at the waypoint denoted by a red square in the airflow map.

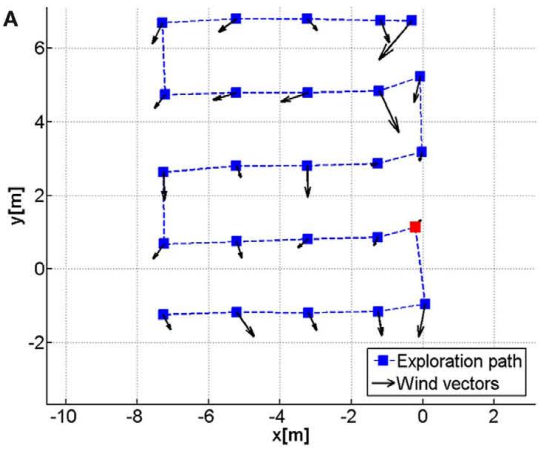

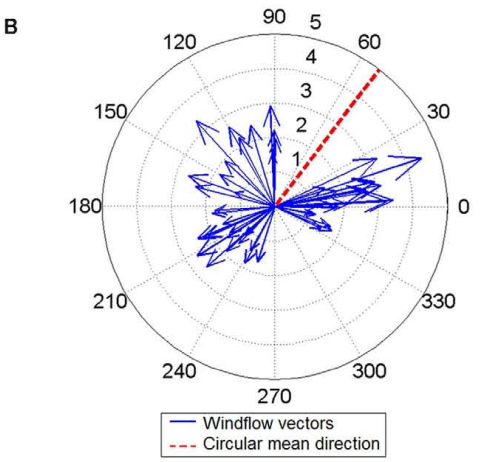

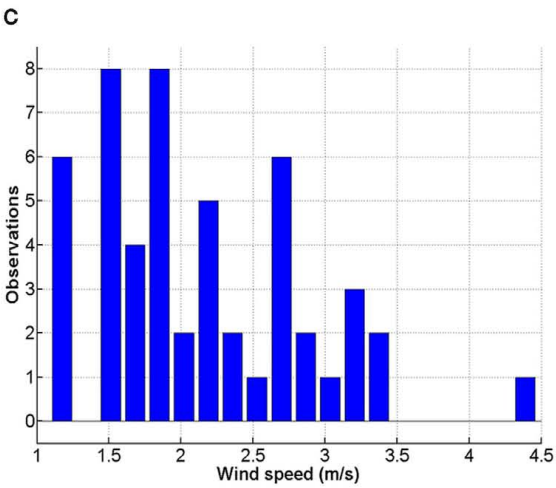

FIGURE 7 | Wind measurements collected during one experimental trial, Örebro University courtyard set-up. (A) Airflow map. The squares denote the points where the robot stopped to collect measurements. The dashed line denotes the robot trajectory. **(B)** Polar plot of the measurements acquired at the waypoint denoted by a red square in the airflow map. **(C)** Wind speed histogram for the measurements acquired at the waypoint denoted by a red square in the airflow map.

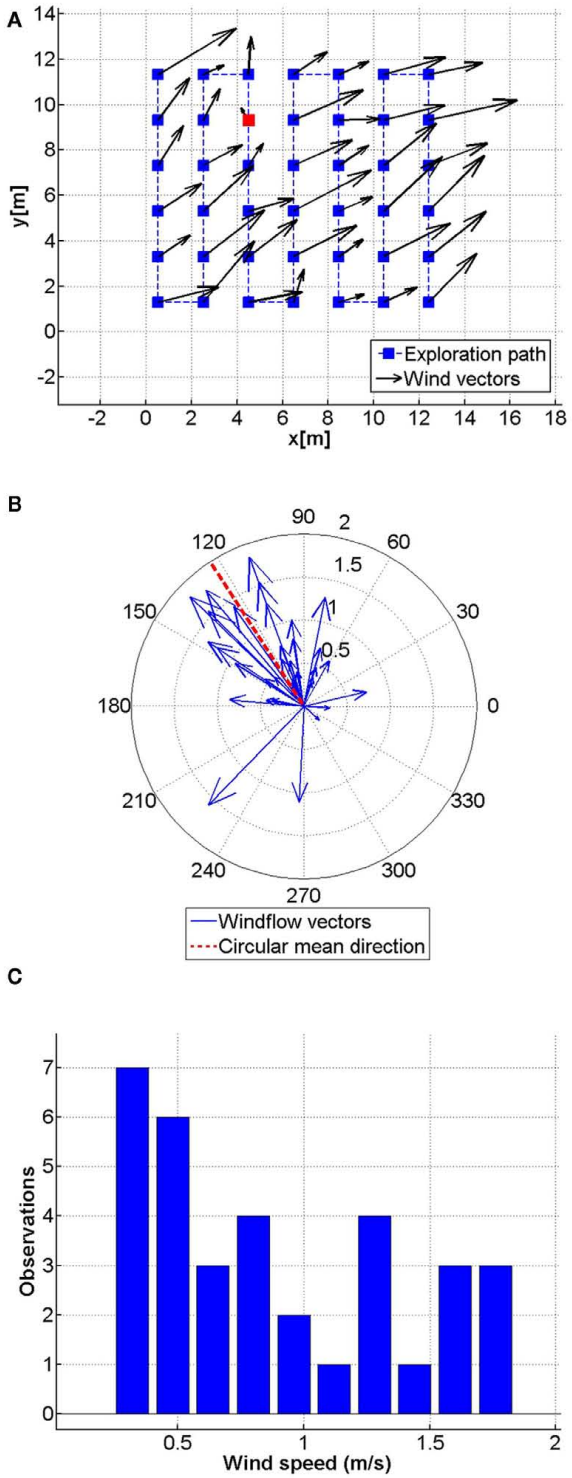

FIGURE 8 | Wind measurements collected during one experimental trial, BAM courtyard set-up. (A) Airflow map. The squares denote the points where the robot stopped to collect measurements. The dashed line denotes the robot trajectory. **(B)** Polar plot of the measurements acquired at the waypoint denoted by a red square in the airflow map. **(C)** Wind speed histogram for the measurements acquired at the waypoint denoted by a red square in the airflow map.

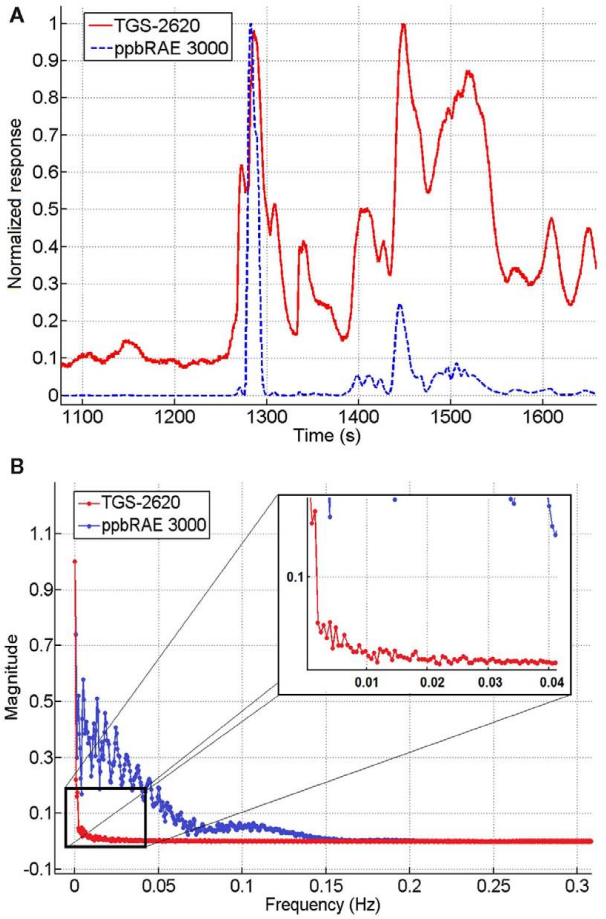

FIGURE 9 | Responses in time domain and frequency domain of a TGS2620 and a PID ppbRAE 3000. The responses were recorded with the robot stopped at a fixed position in the Closed room experimental set-up. **(A)** Time domain responses of the TGS2620 and the PID (ppbRAE 3000) normalized to be between 0 and 1. **(B)** Frequency spectra computed from the response of the TGS2620 and the ppbRAE 3000 in the Closed room experimental set-up.

mean and predictive variance are computed. The model is computed by extrapolating from neighboring measurements weighted by a Gaussian function $\mathcal{N}$ of width σ. Thus, the parameters that regulate the Kernel DM + V algorithm are kernel width σ and cell size *c*. Furthermore, the authors proposed a method to learn the parameters of the algorithm from the measurements by minimizing the average negative log predictive density (NLPD), which is a standard criterion to evaluate distribution models.

For each experimental location, a single sensor was selected according to its sensitivity to generate a gas distribution model. For the three experiments were ethanol was used as the target source, the TGS2620 was selected, while the TGS2611 was used for the experiment conducted with a methane source. The maps were generated using the sensor conductance values recorded while the robot was following the exploration path. The obtained mean and variance maps are shown in **Figure 10**. The learned parameters σ and *c* are listed in **Table 3**. In order to estimate the gas source

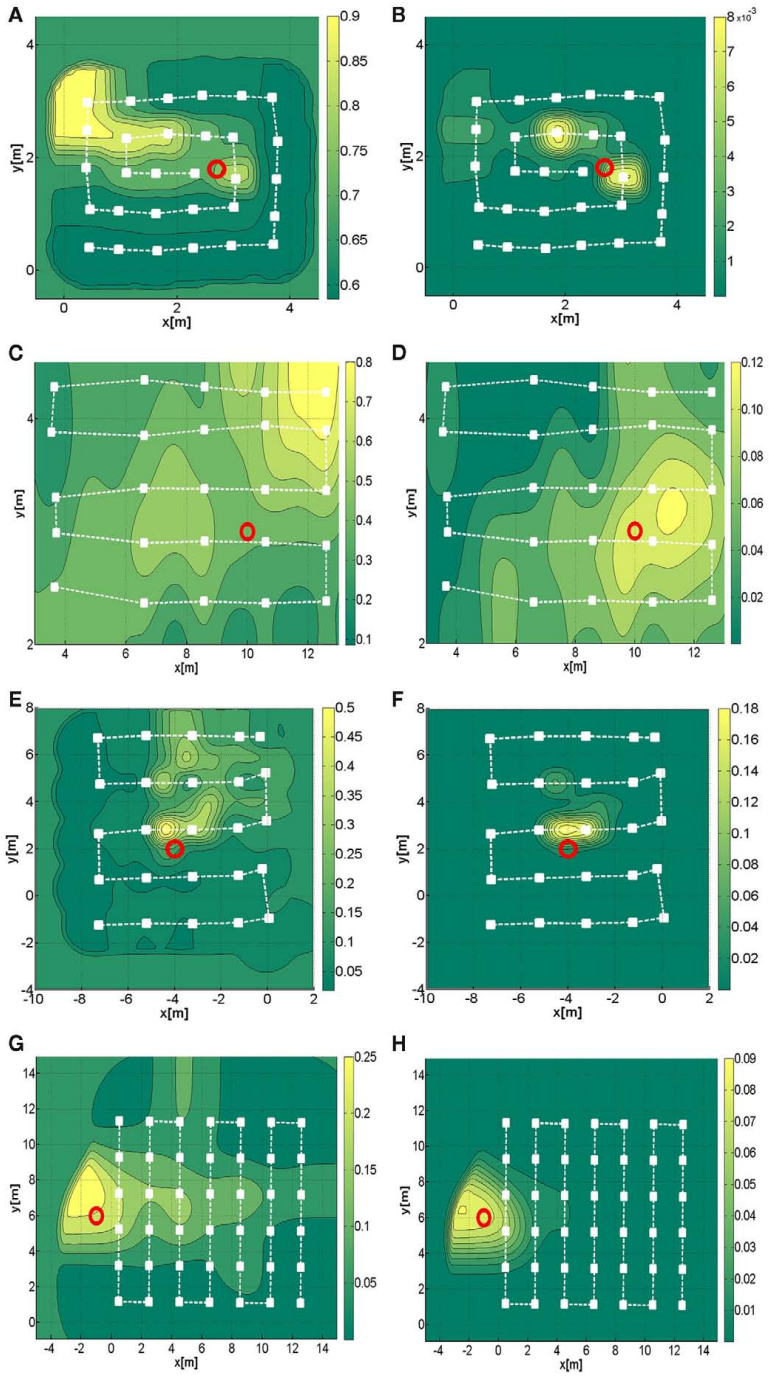

FIGURE 10 | Predictive variance and mean distribution maps obtained with the Kernel DM+V algorithm in the four experimental locations considered in this work. The dashed lines represent the exploration path followed by the robot and the red circle represents the actual gas source location. **(A)** Mean distribution map, closed room. **(B)** Variance distribution map, closed room. **(C)** Mean distribution map, OU corridor. **(D)** Variance distribution map, OU corridor. **(E)** Mean distribution map, OU courtyard **(F)** Variance distribution map, OU courtyard. **(G)** Mean distribution map, BAM courtyard. **(H)** Variance distribution map, BAM courtyard.

locations, we used the variance maps as indicators of the source proximity. It can be noticed that the cells adjacent to the actual source location have a higher variance value, which is represented by brighter color shades in the figures.

4. DISCUSSION

Probably one of the most crucial aspects of research in mobile robot olfaction is the design of the experiments that enable to study and develop systems for airborne chemical monitoring. A major

Table 3 | Kernel DM+V learned parameters used in the four test locations. *c* Stands for grid cell size and σ is the kernel width.

Location	*c* (m)	σ (m)
Closed room	0.10	0.20
Örebro University corridor	0.10	0.44
Örebro University courtyard	0.10	0.40
BAM courtyard	0.10	1.2

technical difficulty is that the dispersion of chemicals in natural environments is difficult to observe since most chemicals produce an invisible plume. In a non-artificial environment, the plume evolution is also difficult to predict *a priori* due to the chaotic dispersal of gas. A second major difficulty is that environmental conditions are often very variable and therefore experiments are hard to repeat. Thus it is difficult to obtain ground truth that can be used to validate experimental results. In order to overcome this limitation, experiments are often carried out under controlled conditions that limit variations from the expected ground truth plume behavior and thus increase repeatability of the experiments. On the other hand, it is hard to predict how the results obtained in such experiments extend to uncontrolled environments. It is very important in the mobile robot olfaction domain that great care is used in describing the experimental set-up in which the claimed results are obtained, and it should be avoided to state that results obtained in an environment with, e.g., steady and controlled airflow trivially extend to more general and complex environments.

In this paper we showed results obtained in four experimental scenarios ranging from a closed room to outdoor areas. In the presented experiments different experimental platforms have been used in order to try to minimize the effect of the choice of a specific robotic system on the presented results. Moreover, the different platforms posed different constraints to the sensors that could be mounted. Given the observed irregularity of the airflow, it is hard to imagine that strategies like *casting*, *surge-spiral*, and *surge-cast*, that try to mimic insect behavior that is adaptive to the up-stream direction, can work reliably. While these strategies have obtained excellent results in a laminar airflow setting, the extension to a turbulent environment is not clear (Lochmatter and Martinoli, 2009). In an attempt to identify possible reasons why the clearly successful reactive tracking behavior of animals cannot be mimicked by current mobile robots, we make two observations: the chemosensing mechanisms currently available are much slower than the biological receptors and mobile robots do not have the same maneuvering capabilities of animals.

According to Justus et al. (2005) the filtering applied by moth antennae is a linear, noise-free representation of odorant concentration in the range of 1–10 Hz, while the gain is reduced for frequencies below 1 Hz. It is argued that the most likely cause for this effect is the adaptation of the receptor cells, a common feature of biological sensory receptors that is most often seen as a slowing or cessation of response to a constant stimulus. According to the data we collected in the four different scenarios mentioned in this paper the bandwidth of the signal collected with MOX gas sensors contains frequencies in the range of 0–0.04 Hz while the signal collected with a PID contains frequencies between 0 and 0.015 Hz. It is striking that currently available chemical sensors stop filtering out the signal in the bandwidth that insects actually can perceive and use for tracking an odor plume. On the other hand, chemical sensors capture the signal in a bandwidth that insects filter out through the adaptation process. Therefore, the perception of the chemical stimulus is totally different for insects compared to state-of-the-art gas-sensitive robots.

Regarding the maneuvering capabilities of animals we consider the moth as an example that has been a predominant source of inspiration for the construction of gas source tracking with mobile robots. According to Kuenen and Cardé (1993) moths fly at a speed of roughly 0.5 m/s and with an average turning rate of roughly 3.5 turns/s. Gas-sensitive mobile robots were instead typically operated at speeds between 0.05–0.1 m/s and can perform much less than 1 turn/s. The limitation in linear speed is mostly to avoid spatial averaging over large areas, which occurs because the gas sensors act as a low-pass filter. (Note, however, that spatial averaging might even help to better localize a gas source in a corridor environment; Lilienthal et al., 2001.) On the other hand, the angular speed is mainly limited by the actuation principles. Based on these differences it seems possible that current gas-sensitive mobile robots are just too slow to perform insect-like reactive steering strategies that allow successful plume tracking in a highly dynamic environment with turbulent airflow.

To conclude, the design of gas-sensitive mobile robots should take into account the limitations of the sensors and robotic platforms currently available. This does not imply that the design of olfactory mobile robots can not be biologically inspired but that mobile robots should not try to directly replicate the movements of insects without a deep understanding how the underlying principles depend on the sensing and actuation capabilities of animals. Instead of mimicking "zigzagging" and "casting" paths, for example, the underlying principles, e.g., the importance of mechanisms to recover from situations in which the animal or robot looses contact with the plume, should be applied when developing gas source tracking approaches. A very interesting example in this regard is the *infotaxis* algorithm which minimizes an entropy function but is nevertheless observed to produce "zigzagging" and "casting" paths similar to those observed in the flight of moths and other animals. Here the underlying principle is probably that the information gain is highest close to the boundary of the plume and this principle could also be used in approaches that do not attempt to zigzag toward a gas source.

In addition to non-biological algorithms that take inspiration from the successful principles reflected in animal behavior, the mobile robot olfaction community should also consider emerging gas sensing technologies like methane laser sensors or infrared cameras that can provide valuable inputs for locations that are meters away from the actual position of the robot.

REFERENCES

Bogner, J., Ahmed, M. A., Diaz, C., Faaij, A., Gao, Q., Hashimoto, S., Mareckova, K., Pipatti, R., and Zhang, T. (2007). *Climate Change 2007: Mitigation. Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge University Press.

Cardé, R., and Willis, M. (2008). Navigational strategies used by insects to find distant, wind-borne sources of odor. *J. Chem. Ecol.* 34, 854–866.

Cleland, T. A., and Linster, C. (2005). Computation in the olfactory system. *Chem. Senses* 30, 801–813.

Gerkey, B., Vaughan, R. T., and Howard, A. (2003). "The player/stage project: tools for multi-robot and distributed sensor systems," in *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR)*, Coimbra, 317–323.

Grasso, F. W., Basil, J. A., and Atema, J. (1998). "Toward the convergence: robot and lobster perspectives of tracking odors to their source in the turbulent marine environment," in *Intelligent Control (ISIC), 1998. Held Jointly with IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA), Intelligent Systems and Semiotics (ISAS), Proceedings*, Gaithersburg, 259–264.

Ihokura, K., and Watson, J. (1994). *The Stannic Oxide Gas Sensor*. Boca Raton: CRC Press.

Ishida, H., Nakamoto, T., Moriizumi, T., Kikas, T., and Janata, J. (2001). Plume-tracking robots: a new application of chemical sensors. *Biol. Bull.* 200, 222–226.

Janata, J. (2009). *Principles of Chemical Sensors*. New York, NY: Springer.

Justus, K. A., Cardé, R. T., and French, A. S. (2005). Dynamic properties of antennal responses to pheromone in two moth species. *J. Neurophysiol.* 93, 2233–2239.

Kowadlo, G., and Russell, R. A. (2008). Robot odor localization: a taxonomy and survey. *Int. J. Rob. Res.* 27, 869–894.

Kuenen, L. P. S., and Cardé, R. T. (1993). Effects of moth size on velocity and steering during upwind flight toward a sex pheromone source by *Lymantria dispar*; (lepidoptera lymantriidae). *J. Insect Behav.* 6, 177–193.

Kuwana, Y., Nagasawa, S., Shimoyama, I., and Kanzaki, R. (1999). Synthesis of the pheromone-oriented behaviour of silkworm moths by a mobile robot with moth antennae as pheromone sensors. *Biosens. Bioelectron.* 14, 195–202.

Lilienthal, A. J., and Duckett, T. (2003). "Experimental analysis of smelling braitenberg vehicles," in *Proceedings of the IEEE International Conference on Advanced Robotics*, Coimbra, 375–380.

Lilienthal, A. J., and Duckett, T. (2004). Building gas concentration gridmaps with a mobile robot. *Robot. Auton. Syst.* 48, 3–16.

Lilienthal, A. J., Duckett, T., Werner, F., and Ishida, H. (2006a). "Indicators of gas source proximity using metal oxide sensors in a turbulent environment," in *Proceedings of the IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, Pisa.

Lilienthal, A. J., Loutfi, A., and Duckett, T. (2006b). Airborne chemical sensing with mobile robots. *Sensors* 6, 1616–1678.

Lilienthal, A. J., Reggente, M., Trincavelli, M., Blanco, J. L., and Gonzalez, J. (2009). "A statistical approach to gas distribution modelling with mobile robots – the kernel dm+v algorithm," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, St. Louis.

Lilienthal, A. J., Reiman, D., and Zell, A. (2003). "Gas source tracing with a mobile robot using an adapted moth strategy," in *Proceedings of the Autonome Mobile Systeme, 18. Fachgespräch* (Karlsruhe: GDI), 150–160.

Lilienthal, A. J., Wandel, M. R., Weimar, U., and Zell, A. (2001). "Sensing odour sources in indoor environments without a constant airflow by a mobile robot," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seoul.

Lochmatter, T. (2010). *Bio-Inspired and Probabilistic Algorithms for Distributed Odor Source Localization Using Mobile Robots*. Ph.D. thesis, Lausanne.

Lochmatter, T., and Martinoli, A. (2009). "Tracking odor plumes in a laminar wind field with bio-inspired algorithms," in *11th International Symposium on Experimental Robotics 2008 (ISER 2008)*, Vol. 54 (Athens: Springer), 473–482.

Moraud, E. M., and Martinez, D. (2010). Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Front. Neurorobotics* 4:1. doi:10.3389/fnbot.2010.00001

Neumann, P., Asadi, S., Lilienthal, A. J., Bartholmai, M., and Schiller, J. H. (2012). Micro-drone for wind vector estimation and gas distribution mapping. *IEEE Robot. Autom. Mag.* (to appear).

Neumann, P., Bartholmai, M., Schiller, J. H., Manolov, M., and Wiggerich, B. (2010). "Self optimizing search and characterization of gaseous hazardous substance sources using a micro-drone: a new approach to determine wind speed and direction," in *IEEE International Workshop on Robotic and Sensors Environments (ROSE)*, Phoenix.

Pyk, P., Badia, S. B. I., Bernardet, U., Knüsel, P., Carlsson, M., Gu, J., Chanie, E., Hansson, B., Pearce, T., and Verschure, P. J. (2006). An artificial moth: chemical source localization using a robot based neuronal model of moth optomotor anemotactic search. *Auton. Robots* 20, 197–213.

Reggente, M., and Lilienthal, A. J. (2009). "Using local wind information for gas distribution mapping in outdoor environments with a mobile robot," in *Sensors, 2009 IEEE*, Christchurch, 1715–1720.

Roberts, P., and Webster, D. (2002). *Turbulent Diffusion. Environmental Fluid Mechanics Theories and Application*. Reston, VA: ASCE Press.

Russell, R. A., Bab-Hadiashar, A., Shepherd, R. L., and Wallace, G. (2003). A comparison of reactive robot chemotaxis algorithms. *Rob. Auton. Syst.* 45, 83–97.

Scharff, H. (2008). *Achieving Adequate Control of Landfill Gas in Europe*. Waste Management World. Available at: http://www.waste-management-world.com

Shraiman, B., and Siggia, E. D. (2000). Scalar turbulence. *Nature* 405, 639–646.

Trincavelli, M. (2011). Gas discrimination for mobile robots. *Künstliche Intelligenz* 25, 351–354.

Trincavelli, M., Coradeschi, S., and Loutfi, A. (2009). Odour classification system for continuous monitoring applications. *Sens. Actuat. B Chem.* 139, 265–273.

Vergassola, M., Villermaux, E., and Shraiman, B. I. (2007). 'Infotaxis' as a strategy for searching without gradients. *Nature* 445, 406–409.

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.

Received: 05 October 2011; paper pending published: 07 November 2011; accepted: 17 December 2011; published online: 12 January 2012.

*Citation: Hernandez Bennetts V, Lilienthal AJ, Neumann PP and Trincavelli M (2012) Mobile robots for localizing gas emission sources on landfill sites: is bio-inspiration the way to go? Front. Neuroeng. **4**:20. doi: 10.3389/fneng.2011.00020*

frontiers in
NEUROENGINEERING

ORIGINAL RESEARCH ARTICLE
published: 29 October 2012
doi: 10.3389/fneng.2012.00022

Evolving a neural olfactorimotor system in virtual and real olfactory environments

Paul A. Rhodes* and Todd O. Anderson

Evolved Machines, Inc., Mountain View, CA, USA

Edited by:
Ramon Huerta, University of California, USA

Reviewed by:
Sandro Carrara, Ecole Polytechnique Fédérale de Lausanne, Switzerland
Marco Trincavelli, Örebro University, Sweden

****Correspondence:***
Paul A. Rhodes, Evolved Machines, Inc., 855 Maude Avenue, Mountain View, CA 94043, USA.
e-mail: prhodes@evolvedmachines.com

To provide a platform to enable the study of simulated olfactory circuitry in context, we have integrated a simulated neural olfactorimotor system with a virtual world which simulates both computational fluid dynamics as well as a robotic agent capable of exploring the simulated plumes. A number of the elements which we developed for this purpose have not, to our knowledge, been previously assembled into an integrated system, including: control of a simulated agent by a neural olfactorimotor system; continuous interaction between the simulated robot and the virtual plume; the inclusion of multiple distinct odorant plumes and background odor; the systematic use of artificial evolution driven by olfactorimotor performance (e.g., time to locate a plume source) to specify parameter values; the incorporation of the realities of an imperfect physical robot using a hybrid model where a physical robot encounters a simulated plume. We close by describing ongoing work toward engineering a high dimensional, reversible, low power electronic olfactory sensor which will allow olfactorimotor neural circuitry evolved in the virtual world to control an autonomous olfactory robot in the physical world. The platform described here is intended to better test theories of olfactory circuit function, as well as provide robust odor source localization in realistic environments.

Keywords: computational fluid dynamics, odor sensor, olfactory, robotics, simulation, virtual world

INTRODUCTION

Brain sensory and control systems evolved to enable action which leads to organism survival. The active interplay between sensation, motor action, and the environment is at the heart of the field of evolutionary robotics (Cliff, 1991; Cliff et al., 1993; Nolfi and Floreano, 2000) and has been studied extensively in the context of visuomotor interaction, so-called "active vision" (Churchland et al., 1994; Floreano et al., 2005). However, evolutionary robotics has received limited attention in the context of olfaction, though olfaction is the more evolutionarily primitive sensory system. While olfactory neural circuitry has been profitably studied as a stand-alone sensory system, combining an olfactory sensory representation with an ability to trigger motor activity has not been attempted. Studying the evolution of sensorimotor transformations is particularly apt in olfaction: olfactorimotor function is reflected in the most evolutionarily primitive chemotactic sensory-motor interactions, such as the spin-and-run behavior observed in paramecium where cilia activity is triggered by chemosensitive ion channels (Greenspan, 2007). Further, the olfactory bulb and piriform cortex dominates the cerebrum in the most evolutionarily primitive vertebrates, hagfish and lamprey (Wicht and Northcutt, 1998), reminding us that olfaction was the foundational cortical sense followed much later by the dedicated visual cortex arising in reptiles (Ulinski, 1990). The physical exploration of a plume has been suggested to be critical in nulling background, making sensory-guided motor control directly related to odorant object identification (Best and Wilson, 2004; Rhodes, submitted). The evolutionarily ancient role of olfaction and the role of exploration in olfactory function suggests that the study of olfactorimotor neural systems will identify principles fundamental to sensorimotor neural systems of all modalities.

A number of intriguing studies have recently pointed to the need for simulations of realistic plume environments in which to develop and test models of olfactorimotor control. Trincavelli (2011) emphasized the joint need for an experimental environment reflecting the turbulence and convection characteristic of natural conditions, but in this work imposed a fixed quasi laminar airflow between source and sensor to enable repeated trials under comparable conditions. In an attempt to develop a database for potential use in studies of olfactorimotor control, Ishida's group (Wada et al., 2010) obtained maps of convection and sensor readings at a regular grid of points in both indoor and outdoor environments; however, the difficulty entailed in using physical robotic platforms to simply map turbulent plumes was reflected in the fact that minutes separated each of dozens of serial measurements of the ever-changing convective environment, preventing the ability to use the database to reproduce the sensor and convection experience of a robotic agent in traversing any path other than the one used to collect the data. Recently Bennetts et al. (2012) addressed olfactorimotor control in several natural environments with sensors on board three types of physical robot platform, two wheeled and one aerial, equipped with metal oxide (MOX) and photon ionization detector (PID) sensors. They also quantified the effective sampling rate for these sensors (<0.04 Hz for the MOX and <0.15 Hz for the PID) suggesting that reversibility is far slower

than the 1–10 Hz relevant to biological sensor responses ranging from canine or rat sniff rates to insect olfactory receptor neurons. Trincavelli and colleagues particularly emphasize the joint necessity of olfactorimotor control experiments with rapidly reversible sensors operating in realistic environments, along with the experimental reproducibility required to systematically test models of olfactorimotor control, whether neurally inspired or not. We have set out to develop a simulation system that meets these goals in this work.

With respect to the control algorithms much work has been motivated by the apparent search strategies adopted by animals, and over the last 15 years a compelling body of work on robotic olfactorimotor control has developed (reviewed in Kowadlo and Russell, 2008; McGill and Taylor, 2011). Most of these studies have however incorporated a more traditional robotic sensorimotor controller, and there has thus far been little work embedding a *neural* olfactory sensorimotor system in a simulated (or real) olfactory environment. In this paper we report on the development of a virtual olfactory plume world with a robotic agent controlled by a neural olfactorimotor simulation. We argue for the use of artificial evolution of the parameters controlling the simulated brain, and then describe how the virtual closed world can be linked to a physical robot, both before and after the development of a physical olfactory sensor with the high dimensionality and rapid reversibility needed to enable plume exploration, so that neural olfactorimotor systems can be evolved in virtual and real worlds in tandem. Describing the components we have chosen to assemble for this integrated system, and addressing some of the constraints encountered and the trade-offs entailed, is the purpose of the present work.

MATERIALS AND METHODS

AN ENVIRONMENT FOR THE SIMULATION OF NEURAL SENSORIMOTOR INTERACTION IN A VIRTUAL PLUME WORLD

To explore olfactorimotor circuit function we embedded a simulated sensory-motor system in a virtual world which simulates multiple turbulent plume sources and their interaction with a robot agent in real-time. This enables modeling two effects of motor commands on the sensory experience: movement of the sensor position through the plume, and perturbation of the plume dynamics due to agent movement. The virtual plume environment is crucial for developing artificial olfactorimotor machines: robots that can autonomously locate a plume source in the face of their interaction with the plume during exploration. As reviewed in the Discussion below, there are currently no virtual world/robotic simulation packages which integrate real-time computational fluid dynamics (CFD) solvers capable of incorporating flow fields shaped by the features of the environment (convection sources, temperature differences) in which to situate odorant plumes. We first explored available CFD implementations capable of integration into a virtual world, seeking one computationally light enough to update plumes and convection fields interactively during simulation, in order to allow agent position and velocity to affect plume and convection motion. The following is an outline of the set of elements that we suggest need to be assembled to enable the study of neural olfactorimotor interactions:

1. A computationally efficient fluid dynamics simulator integrated into a virtual world.
2. A model of robotic agent, including positions of the olfactory sensors and motor effectors to be connected to neuronal motor representation.
3. A means of communicating the currently sampled odorant concentration from the robot sensor to the brain simulation, and of communicating the motor unit activity levels, or resulting effector control signals, back to the agent.
4. Simulated source odorants, an odorant background, and a sensor array.
5. A simulated neural system including circuitry for sensory representation, motor units, and the linkage between them.
6. A mapping of motor unit activity to the control of effectors on the agent, so that the firing of motor neurons in the circuit simulation moves the agent in the virtual world.
7. A means to select and optimize simulation parameters to guide the construction of a biologically inspired neural implementation of an olfactorimotor system embedded in such a virtual world. We have adopted the large-scale use of artificial evolution for this purpose, and outline some of the challenges entailed.
8. A means to bring all three simulations (sensorimotor neural, robotic agent, and plume CFD) into correspondence with a physical robot, initially in a hybrid real-virtual environment where the plume simulation generates the sensor signals and the output of the neural simulation drive both a virtual and real robot in tandem.

Moving from simulated olfactorimotor environments into physical agents of course requires the availability of a physical olfactory sensor array. A suitable artificial olfactory sensor must be rapidly reversible with a time constant comparable to that of biological olfactory sensors to extract information about plume spatiotemporal changes during exploration (Wada et al., 2010; Trincavelli, 2011), and high dimensional if it is to represent a wide range of odorants, two fundamental prerequisites not jointly met by existing physical olfactory sensor alternatives. We are engineering such a sensor[1], an array of functionalized carbon nanotube field effect transistors, and briefly reference this work below.

Below, we report the development of a system incorporating the elements enumerated above:

Computational fluid dynamics simulator

Arguably, the most important aspect of a virtual world devoted to olfactory search behavior is the fluid simulator. Typically the robotic system has the goal of discovering the location of an object (often called the "source") by utilizing a stream of sensor signals triggered by the odorant plume emitted by the source. The CFD simulator determines odorant dispersal from the source by constructing a flow field which takes into account multiple physical effects including air entering and exiting the simulation environment (for example, through an open window, under a door-jam, in a heating vent, or out an air return duct), temperature differentials

[1] www.nanosensetech.com

leading to convective currents, and air displacements due to robot movement (detailed further below). The odor plume, represented as either particles or a continuous value in a voxel or other finite element array, is then transported along this flow field. While the simple diffusion of odor molecules is a factor, real plumes of the scale encountered by most laboratory animals are dominated by turbulent flow.

While simulated robotic olfaction plumes are often precomputed or simulated without ongoing interaction with the moving robotic agent in order to reduce simulation overhead (Cabrita et al., 2010), simulations of the interaction of the body of a moving agent of several centimeters in size or larger with plumes from compact objects (for example sources of food) suggest the perturbation of the plume by the moving agent is very significant, and cannot be plausibly neglected, even in first-order approximation (Dickman et al., 2009). Further, many creatures utilize a variety of means of active intake to enhance olfactory function, including the familiar sniffing characteristic of many mammals and the antennule motion drawing in a stream of water crucial for crustacean olfaction. Such active sensing has been used to advantage in robotic olfactory function (e.g., Ohashi et al., 2008). Simulating olfactorimotor behavior which accounts for perturbation of the plume by movement of the agent, as well as incorporates active intake, requires interactive CFD simulation with an update time comparable in speed to the time step update of the sensorimotor control system.

In selecting a CFD implementation we were therefore forced to balance accuracy with the need for interactive plume updates to account for the movement of the robotic agent under constant (and unpredictable) sensorimotor control by the simulated neural control system. We tested a variety of CFD solutions that enabled multiple odors and interaction between flow field and moving objects but with sufficient computational efficiency to allow plume dynamics to be updated in a time comparable to the neuronal simulation update time. While accurate, traditional CFD software packages such as OpenFOAM, and Comsol Multiphysics are designed to simulate steady-state equilibrium dynamics, rather than the transient dynamics we seek to capture here, and further do not run at speeds approaching real-time. Software designed for generating computer graphics, such as Blender and Houdini, offer exceptional rendering of 3D plumes, but focus on the generation of high resolution effects suited for offline dedicated computation without ongoing interaction with a separately simulated robotic controller. We thus found these software packages unsuited for real-time interactive operation.

While exact implementations of the Navier–Stokes equation system must use very short time steps or implicit solvers in order to avoid instability in the system, Jos Stam has invented a method which employs a simplified version of the Navier–Stokes equations optimized for use in interactive applications such as games (Stam, 1999). To solve fluid equations at video rate, Stam crafts a robust version of the equations which are inherently bounded and thus stable, allowing longer time deltas between simulation steps. Though each time step requires solving a few linear systems and thus is expensive computationally, the fact that each step can represent tens of milliseconds means very few steps are required to run at real-time.

Many implementations of the Stam algorithms exist; for the present work we used Java implementations for ease of prototyping. For rendering a rich graphical interface on a desktop computer we used the MSAfluid fluid dynamics library[2] within the Processing graphics framework[3], and for fast computation with no graphics or interface we used an implementation from Karsten Schmidt's Toxiclibs[4]. The libraries facilitate the creation of a grid of cells, each of which contains a flow force vector and a density variable referred to as "dye." The regional neighborhood of flow vectors determines the transfer (or advection) of dye from one cell to another, as well as the effect of flow velocity on the flow field. Dye may be injected (or removed) at any cell, and thus is ideal for representing odorant plumes (see Movie S1 in Supplementary Material). This CFD implementation allows in-simulation manipulation of plume-related variables, such the position and intensity of the simulated plume and wind sources. Crucially, this system enables interaction between robotic agent and plume environment that characterizes the genuinely interactive system we wish to study (see Movie S2 in Supplementary Material). Our Processing source code implementing the Stam algorithm in two dimensions is linked here: www.stanford.edu/∼tanders/CFD

Currently the virtual world, like the CFD, is two dimensional. Three independent plume sources are modeled, representing three distinct odorants rendered in red, green, and blue, respectively (**Figure 1**). Odorant concentration is represented by color intensity at the specified x, y cell coordinate within the scalar field (**Figures 2A–C**). Air movement, modeled as a force vector field, advects the odor plume. A convection source, such as a fan or vent, is incorporated by adding an additional vector to the existing vector in a specified cell or set of cells. The robot body is modeled as a rectangle, with a top-down image of our real robot (a Surveyor SRV-1) mapped on top. The user may reposition and modify the strength of odor and wind sources using a graphical interface. A separate slider-based panel allows adjustment of such global simulation parameters as viscosity, time step delta, odor evaporation, friction between robot and floor, and the magnitude of the robot's effect on the flow vector field (**Figure 3**). In addition, a graph of the concentrations of each odor over time is plotted in separate window (see **Figure 1**, and Movie S3 in Supplementary Material).

Robot agent simulation

In order to interact with the simulated plume, we need to represent the robotic agent within that plume. The agent, whether simulated or physical, will include olfactory and other sensors (e.g., for collision detection), as well as effectors to transport the robot (motor-driven tracks) and provide for active air intake mechanisms ("sniffing" fans). We require a virtual world which captures both the physics of moving robots, such as friction, inertia, collisions with obstacles, as well as the fluid dynamics discussed in the previous section. As the robot carries the virtual sensor through the virtual world, encountering different parts of the plume, the odor concentrations at the sensor fluctuate from moment-to-moment in a realistic manner (**Figure 1**). These concentration values are

[2]http://memo.tv/msafluid_for_processing
[3]http://www.processing.org
[4]http://toxiclibs.org/

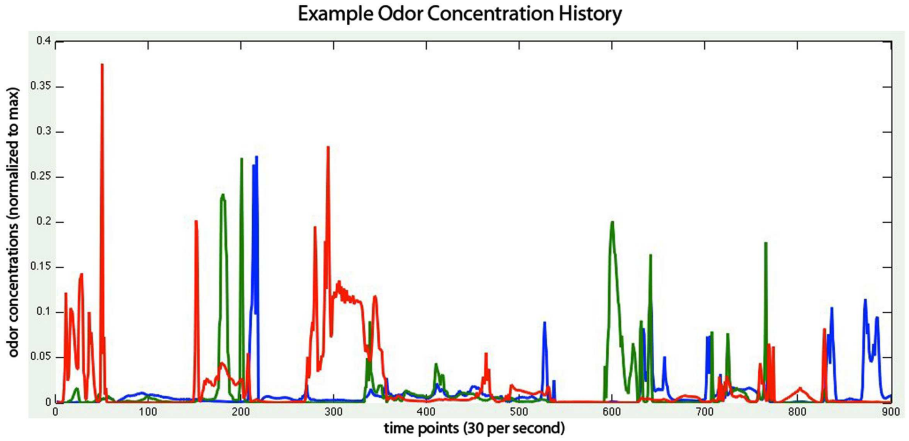

FIGURE 1 | Concentration of three odors (red, green, and blue lines) experienced by the virtual robot as it moves in the virtual environment, with each point representing 1/30 of a second. Thus the *x*-axis reflects 30 s of sensor response. Odor concentration is normalized to the maximum concentration found at the odor source. Note the high rate of change of odorant concentration at the sensor surface, and the independence of relative concentration from the two stationary (blue and red) odorant sources.

sent to the brain (see **Figures 2D–F**), triggering a cascade of activity which eventually activates the motor units. These motor signals are transmitted back to the virtual world to drive the simulated motor effectors used to update the robot position, thus closing the sensorimotor loop via continuous interaction between the virtual environment and separately simulated neural circuitry.

Available virtual worlds designed for simulating robotics are abundant; examples include Microsoft Robotic Simulator (Johns and Taylor, 2008), Webots (Michel, 2004), Player/Stage (Cabrita et al., 2010), and the Robotic Operating System (ROS) promulgated by Willow Garage. Additionally many software packages designed to facilitate creation of video games contain many of the necessary features, including game engines such as Quake (Harvey et al., 2009), Blender (Echeverria et al., 2011), Irrlicht (Ettlin et al., 2005), Unity (Craighead et al., 2007). However, aside from the simplistic particle systems built into many game engines, these existing solutions have not included objects capable of representing odor plumes. One implementation, PlumeSim (Cabrita et al., 2010) built on top of the open source robot simulator Player[5] may enable integration of plumes into a virtual environment but presently lacks support for interaction between the plume and the exploring robots. Because of the plethora of existing libraries for graphical user interface (GUI) control, fluid simulation, and robot communication, and because of the language used to build the CFD model described in Section "CFD Simulator," we wrote a simple custom model robot, incorporating friction via velocity-proportional speed decrement, collision proximity detection, in Java.

Communication between virtual world and neural simulation

To guide the robot agent toward the source of the plume, we integrated the virtual world with a separately simulated neural sensorimotor system based on the Evolved Machines neural simulator (briefly described below in A Simulated Neural System). The simulated agent in the virtual world was connected to the neural simulation ("brain") via a two-way socket protocol based on the Microsoft socket library winsock2 and the Processing library called "Network[6]". Socket connections allow the brain and CFD-virtual world to be modular processes running on separate computers and compatible with any programming language supporting socket communication.

The amount of data to be passed between the neural simulation and the virtual world is extremely light, just a few tens of bytes per timestep. The concentration of each of the three distinct odorants at the position of the sensor is sent from the virtual world to the Neural Simulator. These values activate the sensor array as described in Section "Simulated Virtual Odorants, Background, and the Sensor Array" below, which in turn activates the mitral and cortical arrays, finally activating the motor units (**Figure 4**). The net right-left, forward-backward movement signal resulting from the pattern of activity in the motor unit array, encoded as described in Section "A Mapping of the Motor Units to the Control of Effectors on the Agent" below, is then relayed back to the virtual world so that the agent moves through and perturbs the plume environment. Neural activity passes from the sensor to motor units in five simulation timesteps, creating a degree of propagation-time based latency in the neural system.

The robot interacts with the plumes by adding fluid forces to the vector representing the leading edge of the robot chassis. These forces are equal and opposite to the force measured at each CFD cell along this vector, and as the robot moves so do these vectors. This method generates realistic-looking perturbations of the plume during robot movement (see Movie S2 in Supplementary Material).

The robot movement may be controlled in a variety of ways. A human user can steer the robot, using keyboard controls. The

[5]http://playerstage.sourceforge.net/

[6]http://processing.org/reference/libraries/net/

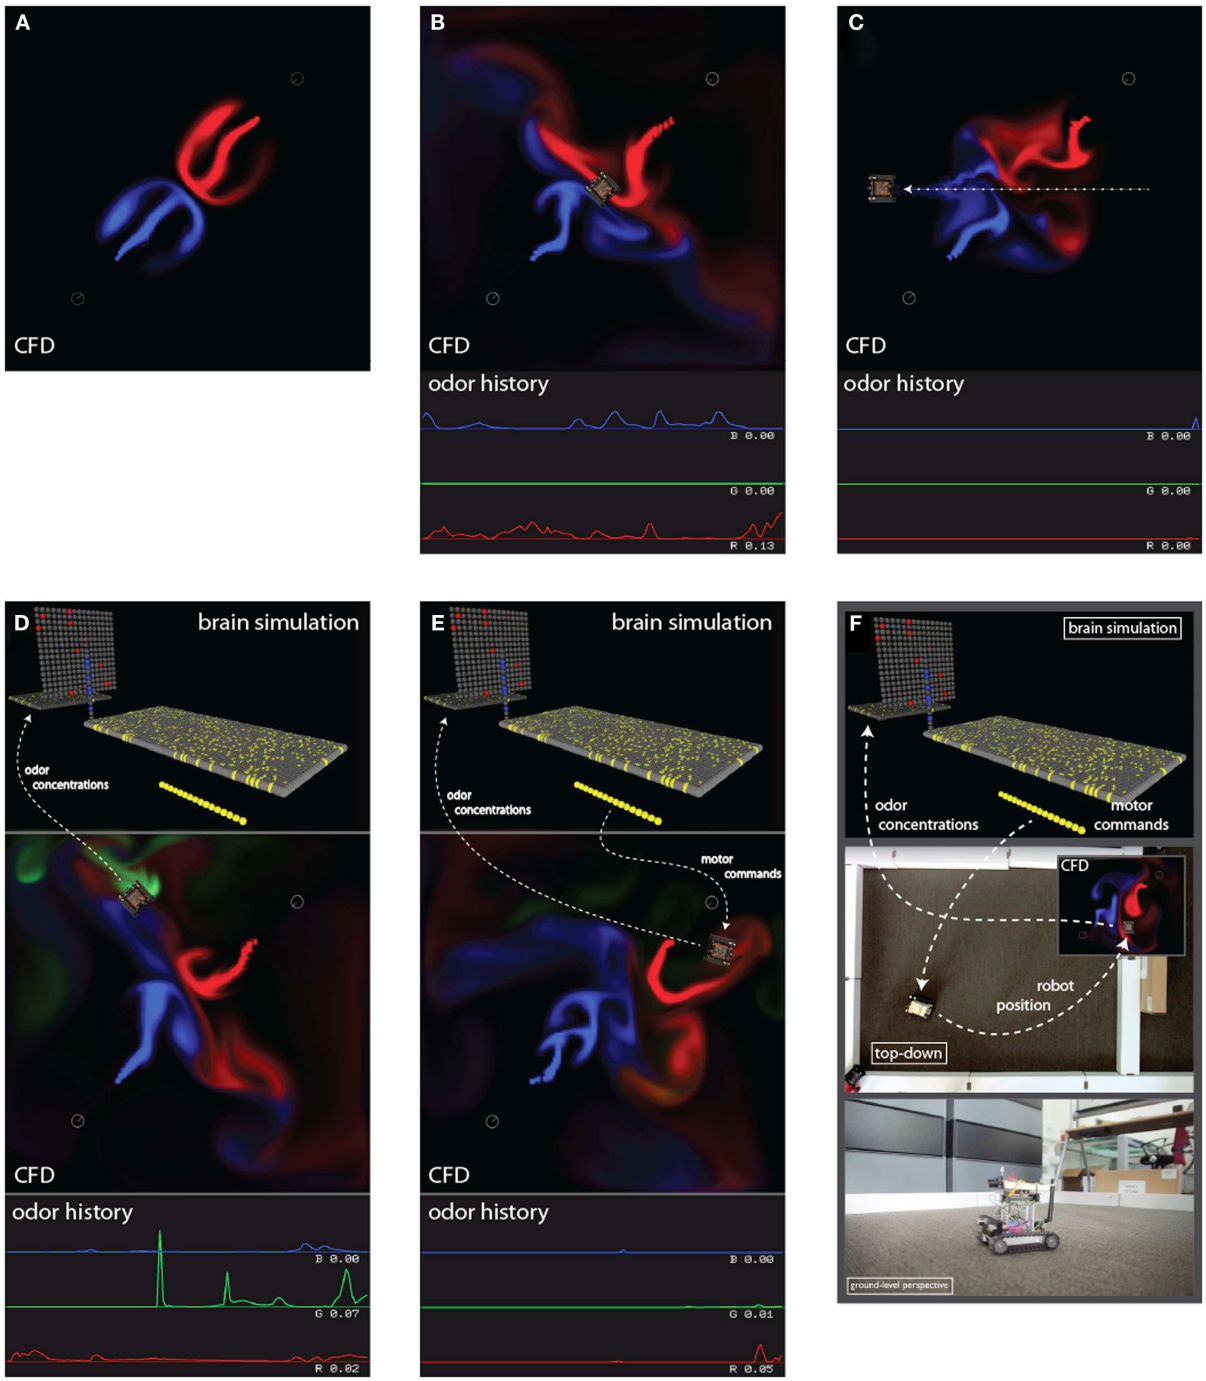

FIGURE 2 | Series of panels show the incremental build up of our simulation framework. (A) CFD alone, near the start of the simulation. **(B)** CFD with robot positioned in the center, and a plot of the odor concentrations experienced by the robot at bottom; the current reading is on the right, and the plot scrolls to the left. **(C)** Robot perturbing the odor plumes; dotted line is the robot path. **(D)** Simulated brain responding to odor concentrations sent by the virtual robot; dashed line shows information flow. **(E)** Simulated brain sending motor commands to the virtual robot. **(F)** Hybrid real-virtual world, where motor commands are sent to a real robot, which is tracked by video to update the position of the virtual robot in the CFD, and therefore send new odor concentrations to the simulated brain.

robot also has a built-in exploratory behavior, which generates motor commands based on the detection of five sensory states. Two of these states are triggered based on absolute odor concentration being either low or high, and two more on rising or falling concentrations relative to a measurement 1 s prior. The fifth behavioral state is a reversing mode, triggered when the virtual robot collides

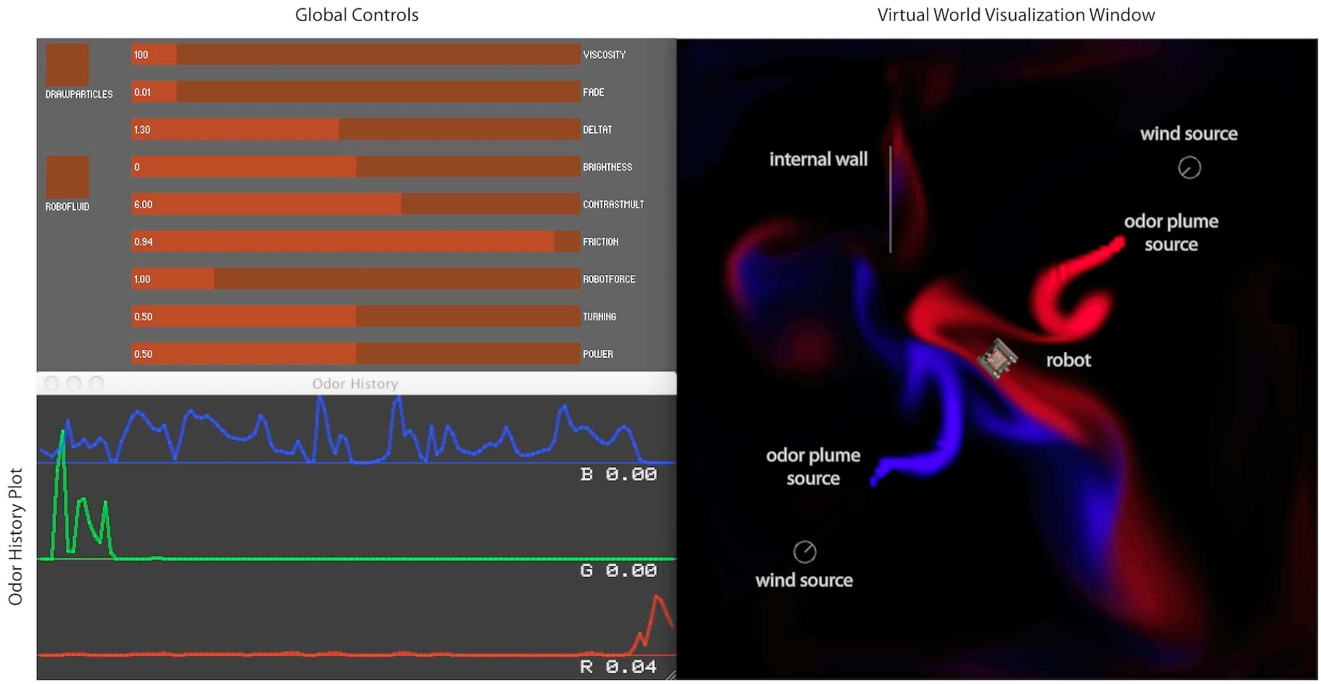

FIGURE 3 | Virtual World graphical user interface. The left side shows a panel containing sliders to control simulation parameters, as well as a scrolling plot of odor concentration history experienced by the robot. The right side shows the virtual world rendering, with two plumes (red and blue), a robot (center), two wind sources (gray circles), and an internal boundary.

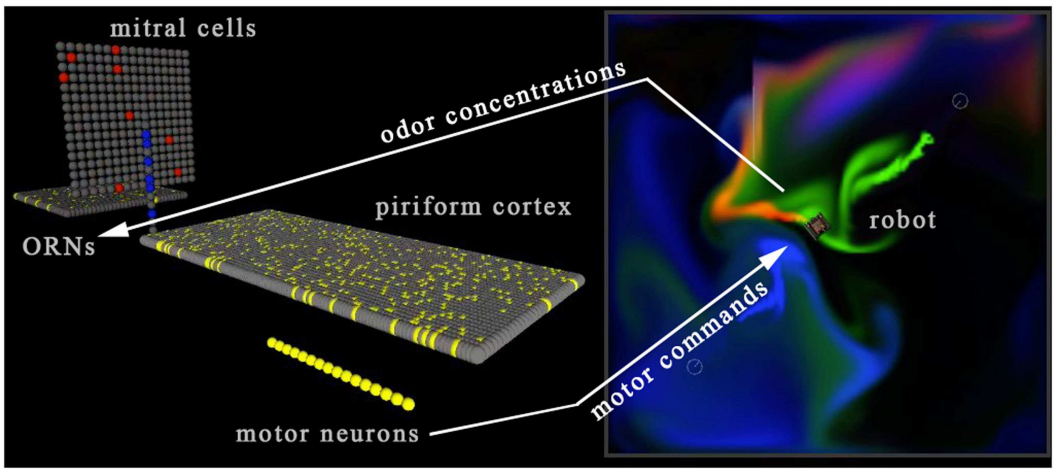

FIGURE 4 | Screenshot of the neural simulator (left side) reacting to odorant concentration from three odorant plumes (red, green, and blue) at the virtual sensor location (right side). A set of 16 motor units activates the forward/backward left/right track treds of the simulated virtual robot, closing the sensorimotor loop with the simulated "brain" receiving sensor information from the virtual world, and via its motor neuron activity sending motor commands into the virtual world. The two otherwise independent software processes were linked by socket-based communications using winsock libraries. See Movie S4 in Supplementary Material.

(reaches zero proximity) with a virtual wall. Finally, the robot can be controlled over the local network, using socket connections. This last method allows communication with human interface devices like joysticks or mobile phones, or more importantly for this project, a remote brain.

Simulated virtual odorants, background, and the sensor array

Odorants. A set of abstract "motifs" was used to specify both odorants and an array of sensors: odorants were defined by the degree (a real number between 0.0 and 1.0) to which they exhibit each of a set of 10 abstract "motifs" (by analogy to the molecular

motifs that correlate with glomerular activation; c.f. Mori et al., 2006). Two odorants could exhibit the same two motifs, though in different relative and absolute degree. The concentration of an odorant is multiplied by motif degree to compute the concentration of the motif present. It is the concentration of motif (not odorant *per se*) that activates the sensors. The background odorant environment comprises a set of additional odorants, constant in space and time but for zero-mean stochastic variation.

Sensors. Input to the neural system initiated with the activation of an array of 256 or 1,024 sensors, each defined by their affinity to several of the motifs chosen at random from the 10 available motifs. For the *i*th sensor the affinity associated with motif α is defined as the inverse of the concentration K_i,α of that motif α (not the associated odorant) necessary to elicit 50% maximal sensor activation. With sensors activation level ranging from 0.0 to 1.0, if a given sensor has affinity $1/K_i$,α to motif α, and an odorant present at concentration C possesses only motif α in degree dα, then the concentration of the motif present is $C{\cdot}d_\alpha$, and at odorant concentration $C_{50} \equiv K_i,\alpha/d_\alpha$, sensor i reaches activation level 0.5. Assuming a first-order interaction between sensor and the concentration of motif α and linear activation of the sensor itself, the steady-state activation level A_α of sensor i in the presence of an odorant at concentration C is given by:

$$A_\alpha = \frac{C \cdot d_\alpha}{C \cdot d_\alpha + K_{i,\alpha}}$$

If an odorant exhibits multiple motifs it is assumed that each activates an independent receptor with first-order kinetics which combine linearly and act jointly to activate the sensor; thus in steady-state the activation level of sensor i presented with an odorant possessing multiple motifs is:

$$A_\alpha = \frac{C}{C + \sum_\alpha K_{i,\alpha} \big/ d_\alpha}$$

where the sum runs over all motifs α present in the odorant to which the sensor i also has affinity. If one assumes that instead of instantaneous activation receptors have a finite rate of activation and deactivation toward the equilibrium activation level, then we instead have:

$$\frac{dA_\alpha}{dt} = r_f \cdot \left(\bar{A}_\alpha - A_\alpha\right) - r_\beta \cdot A_\alpha$$

where r_α and r_β are the rates of activation and deactivation of the sensor respectively. In this work, it was assumed that sensors activated rapidly in comparison to the rate of odorant variation, and so instantaneous activation of sensors was adopted.

A simulated neural system

A very brief description of a subset of the Evolved Machines Neural Simulator used for this work follows. The self-organization of wiring during sensory experience (Rhodes and Taba, 2007), for which this system was built, was de-activated during the present study, which was primarily concerned with putting in place the virtual world and software mechanisms to allow interaction with the a neural sensorimotor circuitry.

Piriform cortex. Pyramidal neurons in piriform cortex and Kenyon cells in the insect homolog both receive input on a set of vertically oriented apical dendrites traversed by a horizontal sheaf of afferent axons from the olfactory bulb and antennal lobe respectively. In a typical cortical pyramidal cell these apical dendrites number between 8 and 50, depending upon species, are largely equal in rank, so that given equivalent excitation it is plausible that each branch could make a comparable contribution to somatic depolarization. The Kenyon cells of the Mushroom Body are similar. Inputs from bulb to piriform cortex are very widely distributed (Stettler and Axel, 2009; Nagayama et al., 2010), as is the projection from antennal lobe to Mushroom Body (Jortner et al., 2007), reflecting a remarkable similarity in architecture between insect and vertebrate olfactory systems. Motivated by these anatomical observations, a cortical array of 4,096 neurons (versus 50,000 in the locust Mushroom Body), each with eight identical branches was constructed, with each neuron receiving 160 inputs, 16 excitatory and 4 inhibitory on each branch, from the bulb output units. Thus the piriform cortex incorporated 655,360 synapses. Branches were independent thresholded units (Rhodes, 1999; but see Bathellier et al., 2009), a branch-spike based model of integration encouraged by indirect evidence suggesting that regenerative branch-level spikes are produced in these branches *in vivo* in insects (Laurent et al., 1993), with a neuron activated in turn by the firing of a threshold number of its branches. Branch threshold was a global parameter that homeostatically adjusted during calibration periods. A model neuron of this type can be considered a detector of the presence of a member of a family of subset detectors, well suited mathematically to orthogonalize overlapping inputs that may represent different objects (Rhodes, 2008). As *in vivo* assessment of branch electrogenesis in vertebrate olfactory cortex pyramids has not yet been made, alternative integrative models were considered, including linear dendrites each of which conveyed their unthresholded summed input to the soma, and sigmoidally activated dendrites (Poirazi et al., 2003) which transformed the linear sum of their inputs with a sigmoid and conveyed the resulting value to the soma. Thresholded units performed better in concentration-invariant olfactory classification in preliminary studies and so were adopted for this study.

Mitral cells. A highly simplified "mitral cell" layer consisting of 256 units with a single dendritic branch was utilized, simply to receive the output of a cluster of four of the 1,024 sensor neurons (**Figure 5**). The interaction between the granule cells and mitral cells present in the vertebrate bulb was neglected in this work, and as a consequence the use of the term "mitral cell" is made only to signify the position of this second layer of units in the flow of activity from sensors to cortex.

Synapses and synaptic depression. There were five classes of synapses in this simulated system: (1) sensor neuron to mitral cell; (2) mitral cell to feedforward interneuron; (3) mitral cell to the dendrites of cortical pyramidal cells; (4) feedforward interneuron to the dendrites of cortical pyramidal cells; (5) and cortical pyramidal cell to motor neuron. Each was modeled as an additive weight, with duration of excitatory post-synaptic potential (EPSP) or inhibitory post-synaptic potential (IPSP) an evolvable

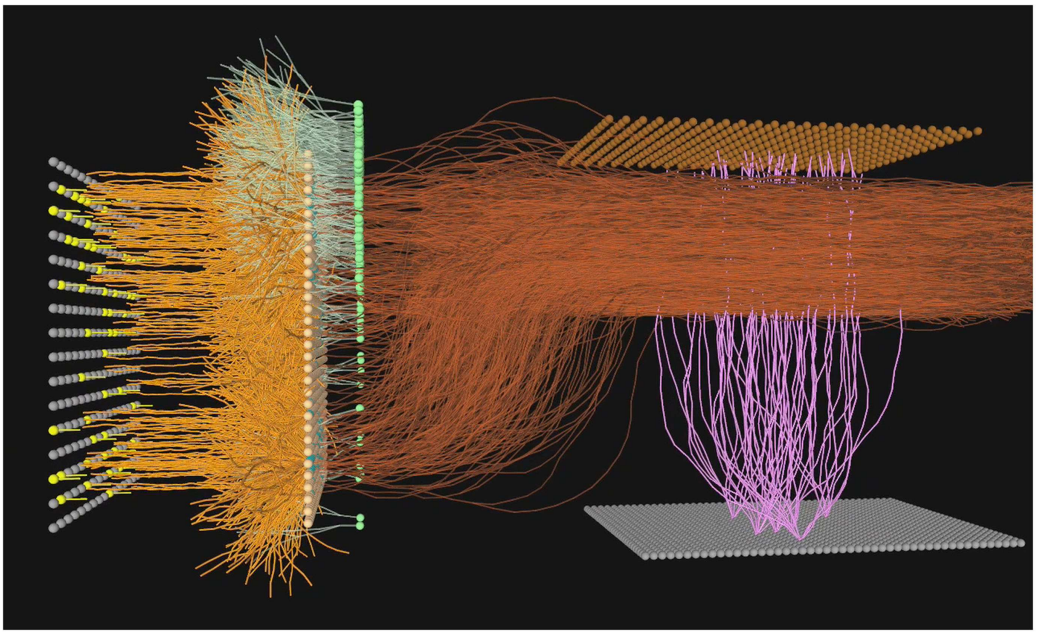

FIGURE 5 | A simulated neural olfactory system, developed using the Evolved Machines neural simulator. Here a sensor sheet (far left) innervates a mitral array (orange dendrites). An inhibitory granule cell layer (green cells and dendrites) is synaptically interconnected with mitral cell lateral dendrites via dendrodendritic reciprocal connections. The resulting mitral cell activity is projected (dark orange axons) to the cortex (sheet of gray cells at bottom left; feedforward interneurons shown without dendrites, at top right), where the pattern of activated cortical cells (lavender) represents the sensor information at this timestep. In the first simulations of the sensorimotor system and plume world, the granule cell dendrodendritic interaction was omitted for simplicity, and a set of 16 motor neurons was added, as described in Section "A Mapping of the Motor Units to the Control of Effectors on the Agent" below and illustrated in **Figures 2D–F**.

parameter for each class. Short term synaptic dynamics (depression and facilitation) is present at all cortical synapses (Thomson, 2000), and accordingly here first-order synaptic depression and recover was incorporated at all synapses in the system. During the evolution of system depression and recovery rates for each of the five classes of synapses were among the variables free to evolve to maximize system performance (see below).

Local homeostatic adjustment of excitability. It has become increasingly clear that the local homeostatic adjustment of dendritic and neuronal excitability and of synaptic efficacy, via a myriad of co-active mechanisms, is ubiquitous in both invertebrate and vertebrate neural circuitry (Davis, 2006). We have found that when simulating the activity-dependent wiring of developing neural circuits homeostatic adjustment of excitability is indispensable, because the target neurons must have a useful dynamic range early in the wiring process, where there are very few connections per neuron so that excitability must be high to enable some postsynaptic activity, as well as and later when hundreds or thousands of connections have been made, so excitability must decline to present constant activation. Here the thresholds of the individual dendrites of cortical pyramids, as well as firing threshold of the pyramidal cells themselves, independently locally self-adjusted to maintain a target average firing rate. The dendritic and somatic average activity rates were independently evolved, while rates of adjustment and the time constant used to compute the activity average were fixed.

A mapping of the motor units to the control of effectors on the agent

While the circuitry and synaptic physiology of the insect and vertebrate olfactory sensory system has been extensively characterized, less information is available, in either insect or vertebrate, regarding the motor system controlled by or affected by olfactory sensory representations. In order to establish an interactive connection between the sensor input from the virtual world to the simulated "brain" and the motor commands sent from the neural system back to move the agent in the virtual world, we connected the cortical array described above to an array of 16 "motor" neurons, each receiving 256 connections from the 4,096 cortical neurons, drawn at random. Given the set of motor effectors on board the simulated agent (forward, backward, left, and right controllers) it was then necessary to define the functional connectivity between the 16 motor units and 4 motor effectors. In other words, when one of these motor neurons fires, what effect occurs at the robot's motor effectors? We chose a mapping that incorporates the "size principle," which in physiology refers to the incremental activation of muscle fibers of increasing power to grade effector force in a useful way, allowing for both fine movement, perhaps appropriate for exploration, and more powerful motor action as might be useful during a surge in the direction of a plume upon its location (Henneman et al., 1965). Motor neurons were divided into two populations, left and right, each of which is further divided into another two populations corresponding to forward and backward control. Within each of these four populations a given neuron drives one wheel of the simulated agent (or a track on the physical

Surveyor robot) left/right, forward/backward, depending on the population identity of that neuron, with force 1, 2, 4, or 8 assigned to the four motor units in each directional pool.

The motor effector system of the agent we chose was intended to mirror the two treads of a bulldozer-like physical robot, the Surveyor SRV-1, that we used in the hybrid system described below. Thus the four motor neuron pools are all that are needed to represent the positive or negative motor current values transmitted to control the two track motors. The two motors require four graded control populations instead of two, because negative values are not represented by spike rates. To solve the problem of representing negative numbers, which are readily interpreted as negative current values to drive the tread motor backward, the neural system encodes negative values with a second set of units for each tread. The activity of each set of motor units is then added together to compute a net current value for the motor for each tread, which is finally the actual 8-bit integer relayed back to the virtual world to control robot motion, as well as to the physical Surveyor robot in the hybrid system described below.

Artificial evolution as a mechanism for refining parameters used in the neural simulation

The sections above outline the components enabling interaction between a simulated robotic agent situated in a virtual olfactory world, with a simulated olfactory sensor activating a neural olfactorimotor simulation, and motor unit activity controlling the actions of the agent. The emphasis has been on achieving an interactive linkage between "brain" activity leading to motor action, and updated sensor stimulation from contact with a plume in the virtual world. The system we describe is clearly highly complex, with a great many parameters associated with the components of the neural mechanisms, including non-linear dendrites, short term synaptic dynamics, and local homeostatic regulation of excitability, all interacting to result in the dynamic relationship between sensor input and motor output. How does one select parameters for this neural system, and make systematically grounded decisions as to which individually well-studied neural mechanisms should be integrated into this system, or excluded in the interest of parsimony and computational efficiency?

We have chosen to approach this problem by defining a performance measure for olfactorimotor behavior of the neurally controlled simulated agent, and then applying large-scale artificial evolution as a means to search the space of parameters as well as neural mechanisms driven by maximization of this fitness. In recent work (Rhodes, submitted) we have evolved simulated olfactory circuitry driven by purely sensory fitness measures, for example the concentration-invariant identification of odorants in stimulus environments incorporating unknown background. In the framework described here, the development of an interactive virtual olfactory world with a robotic agent controlled by the simulated neural olfactorimotor system now brings this work one step closer to true survival-relevant fitness by defining behaviorally defined fitness measures, such as the speed with which the robotic agent achieves a defined proximity to a target odorant source. This measure implicitly combines odorant identification, inherent in reacting to the desired plume amidst other unimportant distracting point sources, and in the midst of distracting background, along with motor control suited to explore and thereby exploit the plume environment to get to the proximity of a source.

Evolution operators, and the selection of parameters subject to evolution. An evolution process entails the choice of the ≈50 parameters for one or more initial "parent" parameter sets, the formation of a first generation from these parents, the specification of a subset of parameters to be subject to variation by mutation, the specification of evolution operators (e.g., mutation, recombination, and selection methods) and their parameters. In each trial, then, a particular parameterization of the neural sensorimotor system controls the action and interaction of the robot in the virtual world, defined by a fixed set of odorant sources and convection boundary conditions. For each such trial fitness (e.g., speed of source localization) is computed and stored, and when a full generation of trials is completed the relative fitness of each individual member of the generation is used to select a set of parents for the next generation. Typically a single generation has minimum on the order of 100 distinct parameter sets ("individuals"), with the selection process iterated for order 100 generations, so that in an evolution process order 10,000–50,000 individual interactive neurally controlled robotic runs are performed. The neural system is simulated entirely on NVIDIA GPU hardware employing their CUDA software framework. The speed of this hardware is such that at the scale of order 10,000 neurons, 100,000 compartments, and several million synapses the neural system side of a simulation of order 10,000 timesteps requires a few minutes, excluding virtual world update. If we consider a model of olfactory function, as suggested by Stopfer et al. (2003), that updates cortical representation in a sequence of cycles clocked by the 20–50 ms beta oscillation for vertebrates and invertebrates respectively, and if we allocate four simulation timesteps with which to update the neural system for each such beta cycle (allowing a single timestep to correspond to a 5–10 ms EPSP), then 10,000 simulation timesteps is approximately 2,500 beta cycles, corresponding to approximately 2 min of sensorimotor exploration of the virtual environment. We chose the CFD simulation with the constraint that the time required for these 2,500 updates of 20–50 ms real-time was also order a few minutes of compute time, so that the neural circuit simulation and CFD simulation with which it interacts runs in comparable times. For the 2-dimensional CFD environment with sources and moving robot, at a 140×140 grid spacing, simulating a 30 ms update requires approximately 2–5 ms on a modern CPU. Thus the 2,500 such updates in a run the interacting CFD requires approximately 10 s While we do not yet know the update time required for the three-dimensional plume world, including perturbation of the convection field by robot motion, it can be 50-fold longer than the interactive CFD update for the 2-dimensional world and remain comparable in time required for update of a neural system of several million synapses. Therefore, depending on the simultaneity of update a 2 min (real-time) robotic exploration of the virtual plume world controlled by a neural sensorimotor system of the scale noted above is computed in approximately 5 min.

The computational resources required, and the use of a farm of GPU's. Given that a parameter evolution process with a 100-member population evolving for 100 generations requires 10,000 such interactive trials, a single GPU-accelerated compute node

would require approximately 1,000 h for a single evolution. If we wish to have these evolutions instead run overnight, say a 12-h period, so each evolution runs results can be analyzed daily, then we can take advantage of the profound parallelism of the evolution process the trials for all members are entirely independent, and so can be run on independent GPU-accelerated nodes, with very data-light communication of fitness results to a control node, and subsequent data-light propagation of the next generation's parameter sets out to the compute nodes, then 80 such GPU-accelerated compute nodes suffices. At Evolved Machines we are completing installation of a 216-node array of NVIDIA GTX-580-accelerated GPU's, which we estimate will allow the evolution of an interactive olfactorimotor system, with fitness computed during several minutes of plume exploration in each simulation, over the course of 150 generations computed for a 150-member population in a 12-h period.

Hybrid virtual world

This paper describes a method for building a robotic odor source localization system which uses a sensorimotor loop modeled after real neural systems. This brain model contains a large number of parameters, and we would like to tune this model to produce behavior which robustly guides the robot to the source the of odor plume. Two facts have lead us to build a virtual world for simulating the interaction between a robot and a set of odor plumes. The first is that we do not yet have a biologically realistic sensor module which can be mounted on a moving robot. The second is that tuning the brain model requires many simulation runs. As compared to running experiments using real robots and real plumes, simulations vastly increase the number of parameter sets which can be tested in a given amount of time.

However, simulations fail to capture all the complexities inherent in real-world robotics and fluids, for example non-homogeneous ground surface topography and friction, sensor noise, wheels/track slippage, uneven motor power. In order to ensure robust behavior, the simulation-tuned brain models must be exposed to these complexities. A hybrid virtual world, in which certain aspects of the real world are simulated, lets us select which of these complexities our system experiences. In addition, the hybrid world lets us get started tuning the neural model even before we have a reliable mobile olfactory sensor. The ultimate goal is to build a real robot to operate in real environments, and we want to be certain our early experiments are grounded in the complexities of the physical world.

Further, wear and tear would render the robot model non-stationary, and while an effective neural controller of physical robots needs to adjust continually for the drift in the physical model (Bongard et al., 2006), just as animal neural circuits do, that would add a challenge to exclusively using physical robots for this study. It is for this reason that we have developed a virtual plume world, where computing power and a farm of compute nodes running in parallel enables running tens of thousands of plume tracking trials in order 10 h instead of order 1000 h.

For these reasons, we sought to increase the realism of our testing system by putting a real robot in the sensorimotor loop. The robot we used is an SRV-1 from Surveyor Corporation[7], which is a small tracked robot, about 10 cm × 12 cm in dimension, and 350 g in its most basic configuration[8]. Motor control commands are relayed to the agent over WiFi (802.11 b/g protocol), as are the signals from the variety of sensors that are conventionally mounted on board, which can include a 1.3 M pixel video camera and infrared proximity sensors. To incorporate the robot into the sensorimotor feedback loop, we need to both control the robot motors using signals from the simulated brain, and send odor concentration signals from the robot to the brain. The first task is relatively straightforward; the socket commands consisting of wheel motor control signals are converted from the activity of the motor units as described in Section "A Mapping of the Motor Units to the Control of Effectors on the Agent" above are converted to Transmission Control Protocol (TCP) commands and passed across the network to a server running on the real robot, which then sends a varying amount of current to the left and right motors. The result is that the simulated brain, activated by virtual sensor activity in the virtual world due to contact between an odorant plume and the location of the virtual sensor, result in activation of motor units in the simulated brain which are converted into motor signals which move the physical Surveyor robot on the floor of the lab (see Movie S5 in Supplementary Material).

Since we do not yet have a physical olfactory sensor for the Surveyor robot, we cannot yet have the sensor signals to the simulated brain come from the physical robot. To bridge the gap we have developed a hybrid system, where the new position of the physical robot, controlled by the simulated motor units as just described in this Section and Section "A Mapping of the Motor Units to the Control of Effectors on the Agent" above, is detected and relayed to the virtual world, where it is used to update the position of the virtual robot in the plume world simulation. This requires tracking the real robot in X and Y, as well as determining the heading (for the purposes of sensor localization). To capture robot posture, we use the infrared blob tracking camera on a WiiMote (Nintendo Corporation) to localize three infrared light-emitting diodes (LEDs) we positioned in an isosceles triangle arrangement on top of the robot (**Figure 6**). The asymmetry of the LEDs lets us determine position and heading from single time points sent over Bluetooth by the WiiMote at 100 Hz, using custom software built with Processing (see text foot note 3) and DarwiinRemote[9]. As the real robot moves, driving by motor commands from the simulated brain, the tracked position and heading are fed into the virtual world to update the posture of the virtual robot. The odor concentrations sent back to the simulated brain are drawn from the robot's position in the simulated plumes, which are perturbed by the motion of the real robot. Thus the virtual robot is no longer controlled by the neural motor units, but rather mirrors the movement of the physical robot, which is tracked as just described. In this way, we are equipped to begin to study the control of a physical robot by a simulated neural sensorimotor systems activated by interaction between a moving olfactory sensor and a simulated plume, and begin to deal with the attendant irregularities (and non-stationarities) including motor response, floor surface traction, power source variability.

[7]http://surveyor.com/

[8]http://www.surveyor.com/SRV_info.html

[9]www.sourceforge.net/projects/darwiin-remote

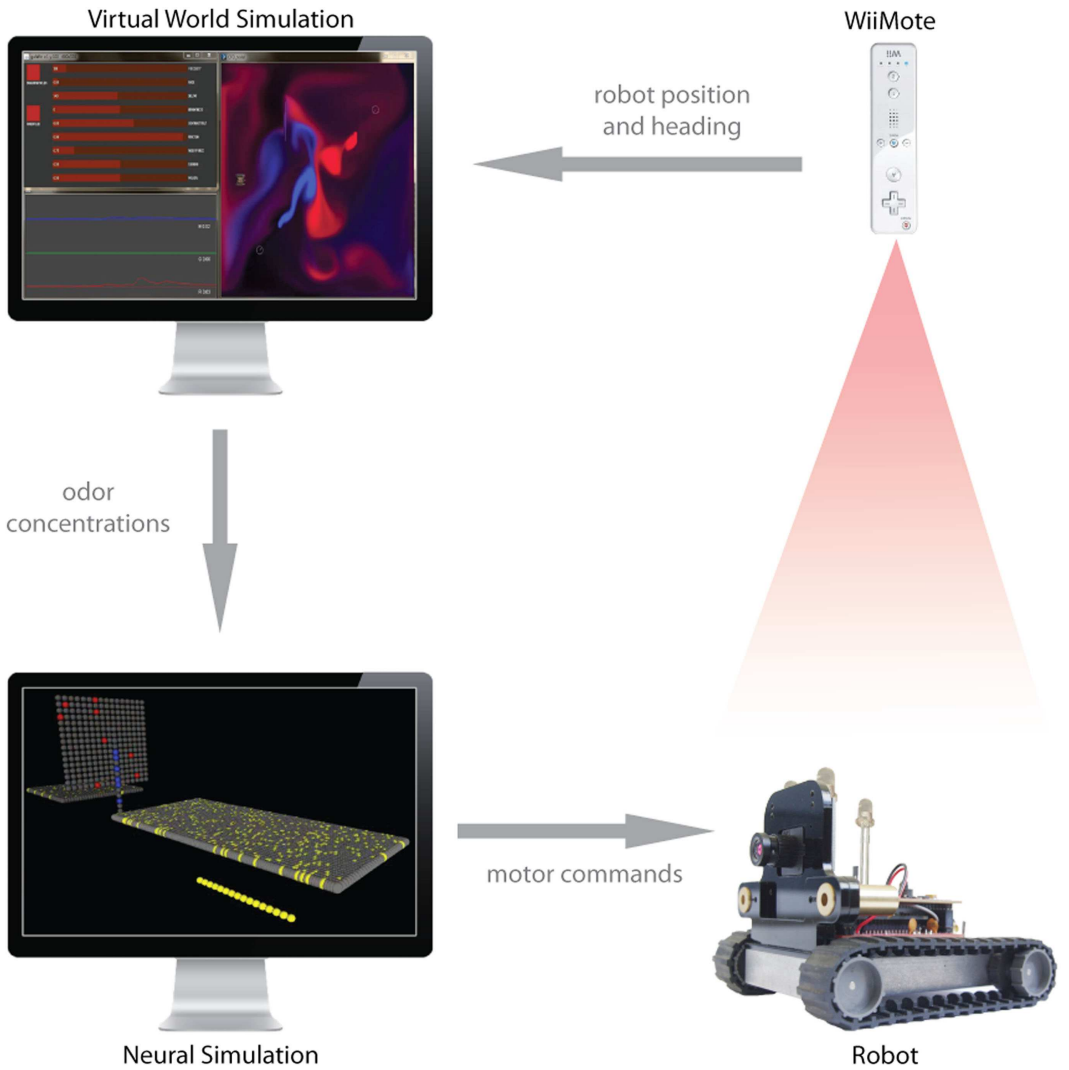

FIGURE 6 | The hybrid virtual-real world experimental setup. Upon receiving a motor command, the physical robot (lower right) executes a maneuver beneath the tracking camera (WiiMote, upper right). Based on isolating the coordinates of three LEDs positioned on the robot's top surface, the coordinates and heading of the robot are sent to the virtual world CFD simulation. The odor concentrations extracted from the robot's position in the CFD are sent to the computer in the lower left, which is running the brain simulation. The odor concentrations are transduced by olfactory receptor neurons, which spread activation through the neural simulation. When the motor neurons become active, a motor command is sent to the physical robot, which completes the feedback loop.

Thus we are able to run 10,000 trials to evolve the parameters of the simulated sensorimotor to optimize as an intermediate step until such time as physical sensors with adequate properties can be engineered and made available, as further described in the next section.

The development of a high dimensional, rapidly reversible, compact, low power artificial olfactory sensor using functionalized nanotube FET arrays

In the foregoing we have outlined a system to study the function of olfactory sensory and sensorimotor circuitry with the assistance of a simulated agent embedded in a virtual world, controlled by the simulated neural "brain." While we expect that the exploration of mechanisms, neural architectures, and parameters choices over order 10,000 trials will continue to require acceleration using a virtual world environments, as noted in the previous section, the olfactory conditions of the real world, and the imperfections of the motor model and effector action (e.g., slippage on the carpet of a wheel) call for transition to a physical machine. To build autonomous olfactory robots driven by reverse-engineered biological neural circuitry requires an olfactory sensor with several crucial properties: (1) High dimensionality. Odorants and odorant mixtures activate biological olfactory sensor arrays identified by expression of a large set of distinct olfactory receptors, ranging from 50 to 150 in the antenna of insects to 900 in the olfactory receptor neuron sheet of olfaction-oriented vertebrates such as canines and rats. The result is a high dimensional dimensional distributed representation, enabling the identification of thousands

of objects and environments. (2) Rapid reversibility, so that a time series of odorant measurements can be collected at an adequate rate as an animal explores a plume environment (as emphasized by Bennetts et al., 2012), a dynamic agility which may be central to nulling background and to facilitate odorant identification in complex environments (Best and Wilson, 2004; Rhodes, submitted). The sample rate which has evolved in of highly olfactory animals, around 5 Hz, as a practical target sampling rate. (3) Further to plausibly be mounted on an autonomous physical robot such a sensor would ideally be low power, compact, rugged, and require no consumables. One of us (PAR) is affiliated with a company (see text footnote 1) engineering a sensor comprised of an array of functionalized carbon nanotube field effect transistors which seeks to meet these requirements, and which will be made available to the academic community. With such a sensor, robotic devices could be used in conjunction with the interactive CFD simulations described here to develop and test models of olfactory sensory and sensorimotor circuitry, situating such simulated neural systems in the context for which that circuitry evolved.

DISCUSSION

SUMMARY OF THE PRESENT WORK

In order to study simulations of neural olfactory circuitry in a motor context we have developed an integrated olfactory virtual world incorporating a computationally light CFD simulation of plumes and convection, in which a simulated robotic agent sends odorant concentration values at its location to a separately simulated neural olfactorimotor system. The resulting continuous stream of activation of a simulated sensory array produces activity in the "cortex" of a simulated neural system incorporating approximately 10,000 cells and a million synapses, for now connected to a simple motor array. The units firing in this motor array in turn relay activation signals to a set of motor effectors on the robotic agent, moving it within the virtual world. Its movement there disturbs the plumes and convection fields being simulated in the ongoing computational fluid dynamic simulation, with the resulting update of odorant concentration at the location of the robot sensor being again relayed to the simulated neural system, closing the sensorimotor loop. We optimize neuronal and circuit parameters and refine the choice of neural mechanisms incorporated in this complex interacting system by the systematic use of artificial evolution, driven by fitness measures chosen to reward performance on relevant sensorimotor tasks such a rapidly locating a plume source, amidst other distracting plumes, and unfamiliar background.

THE PROPOSED PARADIGM IN THE CONTEXT OF RECENT WORK IN NEURAL OLFACTORIMOTOR CONTROL

While there is a rich and growing literature exploring olfactorimotor control in both real and simulated environments (reviewed in Kowadlo and Russell, 2008; McGill and Taylor, 2011), the combination of elements listed above, including a turbulent plume simulation perturbed by a robot which is controlled by a realistic neural simulation optimized using artificial evolution, have not yet been integrated. **Table 1** presents a summary of the literature identifying the simulated environment, method of parameter optimization, motor performance measure, and algorithm space explored, level of interaction between robot and plume, and other characteristics of each of 12 studies, reported within the last 5 years, so that the neural olfactorimotor simulation environment developed and advocated here can be put in context of the distinguished body of work that has informed this field.

With respect to CFD implementation, many studies have implemented turbulent plume simulations using the filament technique of Farrell et al. (2002) or more classic Navier–Stokes (Cabrita et al., 2010). While potentially highly accurate, these equation systems require solution times incompatible with interaction between robot movement (or convection source such as fans, and other mechanisms to achieve sniffing) in real-time during the neural simulation (obviating the ability to run the number of simulations required to do systematic parameter optimization) and so to our knowledge none have enabled the alteration of the plume by a moving robot.

Only a few studies have incorporated neural simulations in guiding olfactory behavior (Mathews et al., 2009; Lopez, 2011), using models built with the IQR simulation framework (Bernardet and Verschure, 2010). Though multiple studies have attempted to emulate the general behavior patterns of odor-seeking insects, often moths (Willis, 2008; Ferri et al., 2009; Cabrita et al., 2010; Lopez, 2011), the use of artificial evolution to systematically search the space of olfactorimotor source localization algorithms (neural or otherwise) is, to our knowledge, unique.

Studies which utilize physical robots (Willis, 2008; Ferri et al., 2009; Mathews et al., 2009; Li et al., 2011; Lopez, 2011) sample real turbulent plume conditions where the perturbation of robot motion is accounted-for, of course, but aside from wind tunnels that ensure laminar flow conditions, it is impossible to ensure that the turbulent plume sensed by the agent is the same from one trial to another. Any method of systematic optimization of the parameters of the simulated system entails a large number of repeated trials which must be sufficiently similarity to make the parameter optimization tractable, suggesting that the use of simulated olfactory environments is indispensable in enabling the exploration of parameter space to an extent that use of physical robots preclude. Once the neural parameters, architecture and mechanisms have been explored in a virtual plume world, incorporation into physical robots is of course a necessary transition to make practical use of the system developed; we therefore emphasized an intermediate step wherein from time to time the physical robot was controlled by the neural motor output, with its resulting position imported to the agent in the plume world, enabling continuous refinement of the virtual world motor model to ensure its relevance to developing motor control of the robot available.

The foregoing review of the literature supports the following enumeration of characteristics of the present system that have not been brought together previously:

1. Control of virtual world agent by a *neural* olfactorimotor system, rather than more conventional robotic control system.
2. Interaction between the simulated robot and the virtual plume and convection field, so that movement of the robot perturbs the plume. This interaction also will enable the addition of (and evolution of the optimum characteristics of) active fluid intake mechanisms such as are employed by biological creatures from lobster to canine to help more efficiently find plume sources.

Table 1 | A survey of studies of olfactorimotor control presented within the last 5 years.

Authors and Year	Neural brain	Parameter optimization via artificial evolution	Closed loop	2D/3D	Real or simulated	# odors	Fluid interaction with robot
Rhodes and Anderson (present work)	Yes	Yes	Yes	2D	Both	3	Yes
McGill and Taylor (2011)	No	No	Yes	2D	Both	Up to 3	No
Li et al. (2011)	No	No	Yes	3D	Real	1	Real
Lopez (2011)	Yes	No	No	3D	Real	2	Real
Lu and Luo (2011)	No	No	Yes	2D	Sim	1	No
Cabrita et al. (2010)	No	No	No	3D	Both	1	No
Moraud and Martinez (2010)				2D	Both	1	No
Li et al. (2011)		No		2D	Sim	1	No
Zarzhitsky et al. (2010)	No	No	Yes	2D	Sim	1	No
Ferri et al. (2009)	No	No	Yes	3D	Real	1	Real
Mathews et al. (2009)	Yes	No	Yes	3D	Real	1	Real
Willis (2008)	No	No	Yes	3D	Real	1	Real

In all but one case, aside from the present study, the sensorimotor control system was not neural, though a variety motor strategies observed in nature were implemented in non-neural control systems. None of the simulation studies incorporated a CFD model that enabled the robot to interact with the plume as it moved, a significant limitation. We advocate the use of a simulated environment to enable the systematic optimization of system parameters, a process that requires the repeated run of plume-robot interaction order 10^3–10^5 times, depending on the number of parameters to be jointly optimized.

3. The routine inclusion of multiple distinct odorant plumes and background odor.
4. The systematic and extensive use of artificial evolution driven by olfactorimotor performance (e.g., to minimize the time to locate a plume source) to specify the parameter values for the neural system and to refine the selection of neural mechanisms to incorporate.
5. A hybrid virtual world with the output of the simulated motor units triggering the motion of a physical robot, the resulting position of which updates the location of the agent in the virtual world, as a bridge to incorporate the realities of motor control of an imperfect physical robot. We note that recently Cabrita et al. (2010) presented a hybrid world with a physical agent wherein a Figaro sensor on board moving through a real environment is assessed in parallel with a simulated agent in a virtual world, though the sensory and motor system was not neural, and there was no interaction between the agent and plume.

OUTLOOK FOR THE FUTURE

While we await development of a physical olfactory sensor with the high dimensionality and rapid reversibility necessary to serve as a front end for a real robot, we have developed an intermediate system connecting the purely simulated neural control system and virtual world to the real world. In this hybrid a physical robot, with its motor control imperfections and non-stationarities, is activated by the motor signals from the simulated neural system, with its position on the lab floor monitored and imported to update the position of the simulated agent, which moves through and perturbs the simulated plume. Finally, we describe the development of a new class of high dimensional, rapidly reversible, low power electronic artificial olfactory sensor which, when available, could be the front end for a fully autonomous neural olfactory robotic, with a neural control system co-evolved in virtual and physical environments. When developed, this olfactory sensor platform will be made available to the research community to both explore the development of working olfactory robotic devices and to enable the study of highly neural simulated olfactory circuitry in the sensorimotor context for which it evolved.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at:http://www.frontiersin.org/Neuroengineering/10.3389/fneng.2012.00022/abstract

REFERENCES

Bathellier, B., Margrie, T. W., and Larkum, M. E. (2009). Properties of piriform cortex pyramidal cell dendrite: implications for olfactory circuit design. *J. Neurosci.* 29, 12641–12652.

Bennetts, V. H., Lilienthal, A., Neumann, P. P., and Trincavelli, M. (2012). Mobile robots for localizing gas emission sources on landfill sites: is bio-inspiration the way to go? *Front. Neuroeng.* 4:20. doi:10.3389/fneng.2011.00020

Bernardet, U., and Verschure, P. M. F. J. (2010). iqr: a tool for the construction of multi-level simulations of brain and behaviour. *Neuroinformatics* 8, 113–134.

Best, A. R., and Wilson, D. A. (2004). Coordinate synaptic mechanisms contributing to olfactory cortical adaptation. *J. Neurosci.* 24, 652–660.

Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science* 314, 1118–1121.

Cabrita, G., Sousa, P., and Marques, L. (2010). "Player/Stage simulation of olfactory experiments," in *IEEE Intelligent Robots and Systems Conference*, 1120–1125.

Churchland, P. S., Ramachandran, V. S., and Sejnowski, T. J. (1994). "A critique of pure vision," in *Large Scale Neuronal Theories of the Brain*, eds J. L. Davis and C. Koch (Cambridge: MIT Press), 23–60.

Cliff, D. T. (1991). "Computational neuroethology: a provisional manifesto," in *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, eds J. A. Meyer and S. W. Wilson (Cambridge: MIT Press), 29–39.

Cliff, D. T., Harvey, I., and Husbands, P. (1993). Explorations in evolutionary robotics. *Adapt. Behav.* 2, 73–110.

Craighead, J., Burke, J., and Murphy, R. (2007). Using the unity game engine to develop sarge: a case study. *Computer* 4552, 366–372.

Davis, G. W. (2006). Homeostatic control of neural activity: from phenomenology to molecular design. *Annu. Rev. Neurosci.* 29, 307–323.

Dickman, B. D., Webster, D. R., Page, J. L., and Weissburg, M. J. (2009). Three-dimensional odorant concentration measurements around actively tracking blue crabs. *Limnol. Oceanogr. Methods* 7, 96–108.

Echeverria, G., Lassabe, N., Degroote, A., and Lemaignan, S. (2011). "Modular open robots simulation engine: MORSE," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, 46–51.

Ettlin, A., Buchler, P., and Bleuler, H. (2005). "A simulation environment for robot motion planning," in *Proceedings of the Fifth International Workshop on Robot Motion and Control*, Poznan, 277–282.

Farrell, J. A., Murlis, J., Long, X., Li, W., and Cardé, R. T. (2002). Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. *Environ. Fluid Mech. (Dordr.)* 2, 143–169.

Ferri, G., Casellia, E., Mattoli, V., Mondini, A., Mazzolai, B., and Dario, P. (2009). SPIRAL: a novel biologically inspired algorithm for gas/odor source localization in an indoor environment with no strong airflow. *Rob. Auton. Syst.* 57, 393–402.

Floreano, D., Suzuki, M., and Mattussi, C. (2005). Active vision and receptive field development in evolutionary robots. *Evol. Comput.* 13, 527–544.

Greenspan, R. J. (2007). *An Introduction to Nervous Systems*. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press.

Harvey, C. D., Collman, F., Dombeck, D. A., and Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461, 941–946.

Henneman, E., Somjen, G., and Carpenter, D. O. (1965). Functional significance of cell size in spinal motoneurons. *J. Neurophysiol.* 28, 560–580.

Johns, K., and Taylor, T. (2008). *Professional Microsoft Robotics Developer Studio.* Birmingham: Wrox Press Ltd.

Jortner, R. A., Farivar, S. S., and Laurent, G. (2007). A simple connectivity scheme for sparse coding in an olfactory system. *J. Neurosci.* 27, 1659–1669.

Kowadlo, G., and Russell, R. A. (2008). Robot odor localization: a taxonomy and survey. *Int. J. Rob. Res.* 27, 869–894.

Laurent, G., Seymour-Laurent, K. J., and Johnson, K. (1993). Dendritic excitability and a voltage-gated calcium current in locust nonspiking local interneurons. *J. Neurophysiol.* 69, 1484–1498.

Li, J.-G., Meng, Q.-H., Wang, Y., and Zeng, M. (2011). Odor source localization using a mobile robot in outdoor airflow environments with a particle filter algorithm. *Auton. Robots* 30, 281–292.

Lopez, L. L. (2011). Moth-like chemosource localization and classification on an indoor autonomous robot, in *On Biomimetics*, ed. L. D. Pramatarova (Rijeka: InTech), 453–466.

Lu, Q., and Luo, P. (2011). A learning particle swarm optimization algorithm for odor source localization. *Int. J. Autom. Comput.* 8, 371–380.

Mathews, Z., Lechón, M., Calvo, J. M. B., Dhir, A., Duff, A., Bermúdez i Badia, S., and Verschure, P. F. M. J. (2009). Insect-like mapless navigation based on head direction cells and contextual learning using chemo-visual sensors. *Intell. Robots Syst.* 2009, 2243–2250.

McGill, K., and Taylor, S. (2011). Robot algorithms for localization of multiple emission sources. *ACM Comput. Surveys* 43, 15.

Michel, O. (2004). Webots: professional mobile robot simulation. *Int. J. Adv. Robot. Syst.* 1, 39–42.

Moraud, E. M., and Martinez, D. (2010). Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Front. Neurorobot.* 4:1. doi: 10.3389/fnbot.2010.00001

Mori, K., Takahashi, Y. K., Igarashi, K. M., and Yamaguchi, M. (2006). Maps of odorant molecular features in the mammalian olfactory bulb. *Physiol. Rev.* 86, 409–433.

Nagayama, S., Enerva, A., Fletcher, M. L., Masurkar, A. V., Igarashi, K. M., Mori, K., and Chen, W. R. (2010). Differential axonal projection of mitral and tufted cells in the mouse main olfactory system. *Front. Neural Circuits* 4:120. doi:10.3389/fncir.2010.00120

Nolfi, S., and Floreano, D. (2000). *Evolutionary Robotics: Biology, Intelligence, and Technology of Self-organizing Machines.* Cambridge, MA: MIT Press.

Ohashi, M., Minagawa, Y., Myoren, Y., and Ishida, H. (2008). "Crayfish robot employing flow induced by waving to locate a chemical source," in *Seventh International Conference on Machine Learning and Applications*, San Diego.

Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal neuron as two-layer neural network. *Neuron* 37, 989–999.

Rhodes, P. A. (1999). "Functional implications of active currents in the dendrites of pyramidal neurons," in *Cerebral Cortex*, Vol. 13, eds P. Ulinski and E. G. Jones (New York: Plenum), 139–200.

Rhodes, P. A. (2008). Recoding patterns of sensory input: higher-order features and the function of nonlinear dendritic trees. *Neural Comput.* 20, 2000–2036.

Rhodes, P. A., and Taba, B. (2007). Growing synthetic neural circuits: activity-dependent wiring in a voxel space. *Soc. Neur. Abst.* 130.19.

Stam, J. (1999). "Stable fluids," in *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99 (New York: ACM Press), 121–128.

Stettler, D. D., and Axel, R. (2009). Representations of odor in the piriform cortex. *Neuron* 63, 854–864.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.

Thomson, A. M. (2000). Molecular frequency filters at central synapses. *Prog. Neurobiol.* 62, 159–196.

Trincavelli, M. (2011). Gas discrimination for mobile robots. *Künstliche Intelligenz* 25, 351–354.

Ulinski, P. S. (1990). The cerebral cortex of reptiles," in *Cerebral Cortex*, Vol. 8A, eds E. G. Jones and A. Peters (New York: Plenum Press), 139–215.

Wada, Y., Trincavelli, M., Fukazawa, Y., and Ishida, H. (2010). Collecting a database for studying gas distribution mapping and gas source localization with mobile robots. *Proceedings of International Confederation of Architectural Museums (IACM)*, Shanghai.

Wicht, H., and Northcutt, R. G. (1998). Telencephalic connections in the Pacific hagfish (*Eptatretus stouti*), with special reference to the thalamopallial system. *J. Comp. Neurol.* 395, 245–260.

Willis, M. A. (2008). Chemical plume tracking behavior in animals and mobile robots. *J. Inst. Navigation* 55, 1–9.

Zarzhitsky, D. V., Spears, D. F., and Thayer, D. R. (2010). Experimental studies of swarm robotic chemical plume tracing using computations fluid dynamics simulations. *Int. J. Intell. Comput. Cybern.* 3, 631–671.

Conflict of Interest Statement: Paul A. Rhodes has an interest in Evolved Machines, Inc. and Nanosense Inc., which are developing simulations of olfactory neural circuitry and their applications in machines, and olfactory sensory for those machines, respectively.

Received: 05 October 2011; accepted: 04 September 2012; published online: 29 October 2012.

Citation: Rhodes PA and Anderson TO (2012) Evolving a neural olfactorimotor system in virtual and real olfactory environments. Front. Neuroeng. ***5****:22. doi: 10.3389/fneng.2012.00022*