



## OPEN ACCESS

## EDITED BY

Thomas Hartung,  
Johns Hopkins University, United States

## REVIEWED BY

Jize Zhang,  
Hong Kong University of Science and  
Technology, Hong Kong SAR, China  
Massimo Salvi,  
Polytechnic University of Turin, Italy

## \*CORRESPONDENCE

Max Olan Smith  
✉ mxsmith@umich.edu

RECEIVED 29 October 2021

ACCEPTED 30 June 2023

PUBLISHED 20 July 2023

## CITATION

Smith MO, Anthony T and Wellman MP (2023)  
Learning to play against any mixture of  
opponents. *Front. Artif. Intell.* 6:804682.  
doi: 10.3389/frai.2023.804682

## COPYRIGHT

© 2023 Smith, Anthony and Wellman. This is an  
open-access article distributed under the terms  
of the [Creative Commons Attribution License  
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction  
in other forums is permitted, provided the  
original author(s) and the copyright owner(s)  
are credited and that the original publication in  
this journal is cited, in accordance with  
accepted academic practice. No use,  
distribution or reproduction is permitted which  
does not comply with these terms.

# Learning to play against any mixture of opponents

Max Olan Smith<sup>1\*</sup>, Thomas Anthony<sup>2</sup> and Michael P. Wellman<sup>1</sup>

<sup>1</sup>Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, United States,  
<sup>2</sup>DeepMind, London, United Kingdom

Intuitively, experience playing against one mixture of opponents in a given domain should be relevant for a different mixture in the same domain. If the mixture changes, ideally we would not have to train from scratch, but rather could transfer what we have learned to construct a policy to play against the new mixture. We propose a transfer learning method, *Q-Mixing*, that starts by learning *Q*-values against each pure-strategy opponent. Then a *Q*-value for any distribution of opponent strategies is approximated by appropriately averaging the separately learned *Q*-values. From these components, we construct policies against all opponent mixtures without any further training. We empirically validate *Q-Mixing* in two environments: a simple grid-world soccer environment, and a social dilemma game. Our experiments find that *Q-Mixing* can successfully transfer knowledge across any mixture of opponents. Next, we consider the use of observations during play to update the believed distribution of opponents. We introduce an opponent policy classifier—trained reusing *Q*-learning data—and use the classifier results to refine the mixing of *Q*-values. *Q-Mixing* augmented with the opponent policy classifier performs better, with higher variance, than training directly against a mixed-strategy opponent.

## KEYWORDS

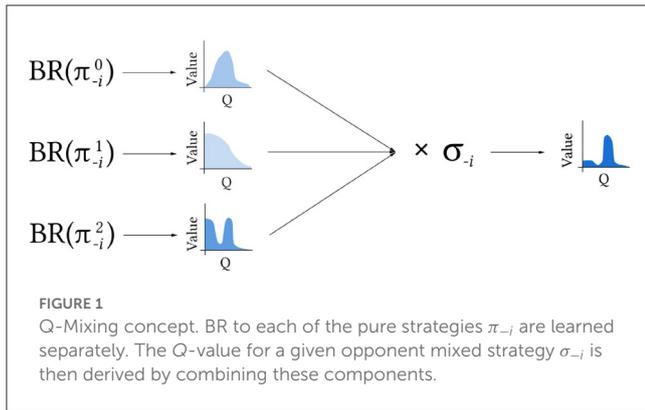
reinforcement learning (RL), transfer learning (TL), deep reinforcement learning (deep RL), value-based reinforcement learning, multi-agent learning

## 1. Introduction

Reinforcement learning (RL) agents commonly interact in environments with other agents, whose behavior may be uncertain. For any particular probabilistic belief over the behavior of another agent (henceforth, *opponent*), we can learn to play with respect to that opponent distribution (henceforth, *mixture*), for example by training in simulation against opponents sampled from the mixture. If the mixture changes, ideally we would not have to train from scratch, but rather could *transfer* what we have learned to construct a policy to play against the new mixture.

Traditional RL algorithms include no mechanism to explicitly prepare for variability in opponent mixtures. Instead, current solutions either learn a new behavior or update a previously learned behavior. In lieu of that, researchers designed methods for learning a single behavior successfully across a set of strategies (Wang and Sandholm, 2003), or quickly adapting in response to new strategies (Jaderberg et al., 2019). In this work, we explicitly tackle the unique problem of responding to new opponent mixtures without requiring further simulations for learning.

We propose a new algorithm, *Q-Mixing*, that effectively transfers learning across opponent mixtures. The algorithm is designed within a population-based learning regime (Lanctot et al., 2017; Jaderberg et al., 2019), where training is conducted against a known distribution of opponent policies. *Q-Mixing* initially learns value-based best responses (BR), represented as *Q*-functions, with respect to each of the opponent's pure strategies. It then transfers this knowledge against any given opponent mixture by weighting the BR



Q-functions according to the probability that the opponent plays each of their pure strategies. The idea is illustrated in Figure 1. This calculation is an approximation of the Q-values against the mixed-strategy opponent, with error due to misrepresenting the future belief of the opponent by the current belief. The result is an approximate BR policy against any mixture, constructed without additional training.

This situation is common in game-theoretic approaches to multiagent RL such as Policy-Space Response Oracles (PSRO; Lanctot et al., 2017) where a population of agents is trained and reused. PSRO iteratively solves a game by adding policies to each agent's strategy set that are best-responses to a mixture of the opponent's current strategy set. Knowledge of the opponent's mixture is available during training, but is currently unused. Smith et al. (2021) showed that PSRO's best-response calculation can use knowledge of the opponent's mixture in combination with Q-Mixing to reduce PSRO's cumulative training time.

We experimentally validate our algorithm on: (1) a simple grid-world soccer game, and (2) a sequential social dilemma game. Our experiments show that Q-Mixing is effective in transferring learned responses across opponent mixtures. We also address two potential issues with combining Q-functions according to a given mixture. The first is that the agent receives information *during play* that provides evidence about the identity of the opponent's pure strategy. We address this by introducing an opponent classifier to predict which opponent we are playing for a particular observation and reweighting the Q-values to focus on the likely opponents. The second issue is that the complexity of the policy produced by Q-Mixing grows linearly with the support of the opponent's mixture. This can make simulation and decision making computationally expensive. We propose and show that policy distillation (Rusu et al., 2015) can compress a Q-Mixing policy while maintaining performance.

#### Key contributions:

1. Theoretically relate (in an idealized setting) the Q-value for an opponent mixture to Q-values for mixture components.
2. A new transfer learning algorithm, *Q-Mixing*, that uses this relationship to construct approximate policies against any given opponent mixture, without additional training.

3. Augmenting this algorithm with runtime opponent classification, to account for observations that may inform predictions of opponent policy during play.
4. Demonstrate that policy distillation can reduce Q-Mixing's complexity (in both memory and computation).

## 2. Preliminaries

At any time  $t \in \mathcal{T}$ , an agent receives the environment's state  $s^t \in \mathcal{S}$ , or an observation  $o^t \in \mathcal{O}$ , a partial state. Even if an environment's state is fully observable, the inclusion of other agents with uncertain behavior makes the overall system partially observable. From said observation, the agent takes an action  $a^t \in \mathcal{A}$  receiving a reward  $r^t \in \mathbb{R}$ . The agent's *policy* describes its behavior given each observation  $\pi : \mathcal{O} \rightarrow \Delta(\mathcal{A})$ . Actions are received by the environment, and a next observation is determined following the environment's transition dynamics  $p : \mathcal{O} \times \mathcal{A} \rightarrow \mathcal{O}$ .

The agent's goal is to maximize its reward over time; called the *return*:  $G^t = \sum_{l=0}^{\infty} \gamma^l r^{t+l}$ , where  $\gamma$  is the discount factor weighting the importance of immediate rewards. Return is used to define the state- and action-values respectively:

$$V(o^t) = \mathbb{E}_{\pi} \left[ \sum_{l=0}^{\infty} \gamma^l r(o^{t+l}, a^{t+l}) \right],$$

$$Q(o^t, a^t) = r(o^t, a^t) + \gamma \mathbb{E}_{o^{t+1} \in \mathcal{O}} [V(o^{t+1})].$$

For multiagent settings, we index the agents and distinguish the agent-specific components with subscripts. Agent  $i$ 's policy is  $\pi_i : \mathcal{O}_i \rightarrow \Delta(\mathcal{A}_i)$ , and the opponent's policy<sup>1</sup> is the negated index,  $\pi_{-i} : \mathcal{O}_{-i} \rightarrow \Delta(\mathcal{A}_{-i})$ . Boldface elements are joint across the agents (e.g., joint-action  $\mathbf{a}$ ).

Agent  $i$  has a strategy set  $\Pi_i$  comprising the possible policies it can employ. Agent  $i$  may choose a single policy from  $\Pi_i$  to play as a *pure strategy*, or randomize with a *mixed strategy*  $\sigma_i \in \Delta(\Pi_i)$ . Note that the pure strategies  $\pi_i$  may themselves choose actions stochastically. For a mixed strategy, we denote the probability the agent plays a particular policy  $\pi_i$  as  $\sigma_i(\pi_i)$ . A *best response* (BR) to an opponent's strategy  $\sigma_{-i}$  is a policy with maximum return against  $\sigma_{-i}$  and may not be unique.

Agent  $i$ 's prior belief about its opponent is represented by an opponent mixed-strategy,  $\sigma_{-i}^0 \equiv \sigma_{-i}$ . The opponent plays the mixture by sampling a policy according to  $\sigma_{-i}$  at the start of the episode. Then agents are locked into the sampled policy's behavior for the entire duration of the episode. The agent's updated belief at time  $t$  about the opponent policy faced is denoted  $\sigma_{-i}^t$ .

We introduce the term *Strategy Response Value* (SRV) to refer to the observation-value against an opponent's strategy.

**Definition 1** (Strategic Response Value). *An agent's  $\pi_i$  strategic response value is its expected return given an observation, when*

<sup>1</sup> Our methods are defined here for environments with two agents. Extension to greater numbers while maintaining computational tractability is a topic for future work.

playing  $\pi_i$  against a specified opponent strategy:

$$V_{\pi_i}(o_i^t | \sigma_{-i}^t) = \mathbb{E}_{\sigma_{-i}^t} \left[ \sum_{\mathbf{a}} \pi(\mathbf{a} | o_i^t) \sum_{o'_i, r_i} p(o'_i, r_i | o_i^t, \mathbf{a}) \delta \right]$$

where  $\delta \equiv r_i + \gamma V_{\pi_i}(o'_i | \sigma_{-i}^{t+1})$ . Let the optimal SRV be

$$V_i^*(o_i^t | \sigma_{-i}^t) = \max_{\pi_i} V_{\pi_i}(o_i^t | \sigma_{-i}^t).$$

From the SRV, we define the *Strategic Response Q-Value* (SRQV) for a particular opponent strategy.

**Definition 2** (Strategic Response Q-Value). *An agent's  $\pi_i$  strategic response Q-value is its expected return for an action given an observation, when playing  $\pi_i$  against a specified opponent strategy:*

$$Q_{\pi_i}(o_i^t, a_i^t | \sigma_{-i}^t) = \mathbb{E}_{\sigma_{-i}^t} [r_i^t] + \gamma \mathbb{E}_{\sigma_{-i}^{t+1}} [V_{\pi_i}(o_i^{t+1} | \sigma_{-i}^{t+1})],$$

where  $r_i^t \equiv r_i(o_i^t, a_i^t, \sigma_{-i}^t)$ . Let the optimal SRQV be

$$Q_i^*(o_i^t, a_i^t | \sigma_{-i}^t) = \max_{\pi_i} Q_{\pi_i}(o_i^t, a_i^t | \sigma_{-i}^t).$$

### 3. Q-Mixing

Our goal is to transfer the Q-learning effort across different opponent mixtures. We consider the scenario where we first learn against each of the opponent's pure strategies. From this, we construct a Q-function for a given distribution of opponents from the policies trained against each opponent's pure strategy.

#### 3.1. Single-state setting

Let us first consider a simplified setting with a single state. This is essentially a problem of bandit learning, where our opponent's strategy will set the reward of each arm for an episode. Intuitively, our expected reward against a mixture of opponents is proportional to the payoff against each opponent weighted by their respective likelihood.

As shown in Theorem 3, weighting the component SRQV by the opponent's distribution supports a BR to that mixture. We call this relationship *Q-Mixing-Prior* and define it in Theorem 3.

**Theorem 3** (Single-State Q-Mixing). *Let  $Q_i^*(\cdot | \pi_{-i})$ ,  $\pi_{-i} \in \Pi_{-i}$ , denote the optimal strategic response Q-value against opponent policy  $\pi_{-i}$ . Then for any opponent mixture  $\sigma_{-i} \in \Delta(\Pi_{-i})$ , the optimal strategic response Q-value is given by*

$$Q_i^*(a_i | \sigma_{-i}) = \sum_{\pi_{-i} \in \Pi_{-i}} \sigma_{-i}(\pi_{-i}) Q_i^*(a_i | \pi_{-i}).$$

*Proof:* The definition of Q-value is as follows (Sutton and Barto, 2018):

$$Q_i^*(a_i) = \sum_{r_i} p(r_i | a_i) \cdot r_i.$$

In a multiagent system, the dynamics model  $p$  suppresses the complexity introduced by the other agents. We can unpack the dynamics model to account for the other agents as follows:

$$p(r_i | a_i) = \sum_{\pi_{-i}} \sum_{a_{-i}} \pi_{-i}(a_{-i}) \cdot p(r_i | \mathbf{a}).$$

We can then unpack the strategic response value as follows:

$$Q_i^*(a_i | \pi_{-i}) = \sum_{a_{-i}} \pi_{-i}(a_{-i}) \sum_{r_i} p(r_i | \mathbf{a}) \cdot r_i.$$

Now we can rearrange the expanded Q-value to explicitly account for the opponent's strategy. The independence assumption enables the following re-writing by letting us treat the opponent's mixed strategy as a constant condition:

$$\begin{aligned} Q_i^*(a_i | \sigma_{-i}) &= \sum_{r_i} \sum_{\pi_{-i}} \sigma_{-i}(\pi_{-i}) \sum_{a_{-i}} \pi_{-i}(a_{-i}) p(r_i | \mathbf{a}) \cdot r_i \\ &= \sum_{\pi_{-i}} \sigma_{-i}(\pi_{-i}) \sum_{a_{-i}} \pi_{-i}(a_{-i}) \sum_{r_i} p(r_i | \mathbf{a}) \cdot r_i \\ &= \sum_{\pi_{-i}} \sigma_{-i}(\pi_{-i}) Q_i^*(a_i | \pi_{-i}). \end{aligned}$$

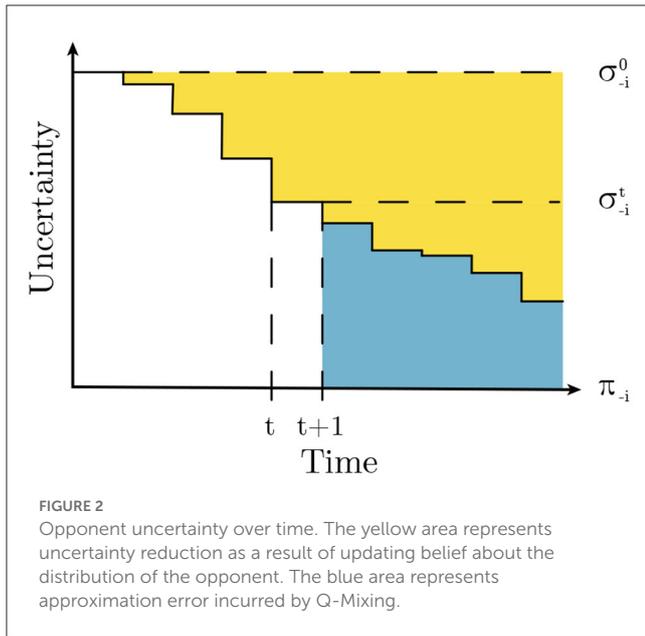
While this setting facilitates exact mixing of Q-values, its applicability to general games is limited. In the following subsections, we will explore how we can trade-off exactness in the mixing of Q-values in order to broaden its use to a larger classes of games.

#### 3.2. Leveraging information from the past

Next, we consider the RL setting where both agents are able to influence an evolving observation distribution. As a result of the joint effect of agents' actions on observations, the agents have an opportunity to gather information about their opponent during an episode. Methods in this setting need to (1) use information from the past to update its belief about the opponent, and (2) grapple with uncertainty about the future. To bring Q-Mixing into this setting we need to quantify the agent's current belief about their opponent and their future uncertainty.

During a run with an opponent's pure strategy drawn from a distribution, the actual observations experienced generally depend on the identity of this pure strategy. Let  $\psi_i: \mathcal{O}_i^{0:t} \rightarrow \Delta(\Pi_{-i})$  represent the agent's current belief about the opponent's policy using the observations during play as evidence. From this prediction, we propose an approximate version of Q-Mixing that accounts for past information. The approximation works by first predicting the relative likelihood of each opponent given the current observation. Then it weights the Q-value-based BRs against each opponent by their relative likelihood.

Figure 2 provides an illustration of the benefits and limitations of this new prediction-based Q-Mixing. At any given timestep  $t$  during the episode, the information available to an agent about the opponents may be insufficient to identify perfectly their policy. In the figure, the yellow area above a timestep represents the



uncertainty reduction from an updated prediction of the opponent  $\sigma_{-i}^t$  compared to the baseline prediction of the prior  $\sigma_{-i}^0$ . Upon perception of the next observation  $t + 1$ , the agent may collect additional evidence informing the identity of the opponent further reducing the uncertainty. Crucially, this definition of Q-Mixing does not consider updating the opponent distribution from new information in the future (blue area in Figure 2).

Let the previously defined  $\psi$  be the *opponent policy classifier* (OPC), which predicts the opponent's policy. In this work, we consider a simplified version of this function, that operates only on the agent's current observation  $\psi_i: \mathcal{O}_i \rightarrow \Delta(\Pi_{-i})$ . We then augment Q-Mixing to weight the importance of each BR as follows:

$$Q_{\pi_i}(o_i, a_i | \sigma_{-i}) = \sum_{\pi_{-i}} \psi_i(\pi_{-i} | o_i, \sigma_{-i}) Q_{\pi_i}(o_i, a_i | \pi_{-i}).$$

We refer to this quantity as *Q-Mixing*, or *Q-Mixing-X*, where X describes  $\psi$  (e.g., *Q-Mixing-Prior* defines  $\psi$  as the prior uncertainty of the opponent  $\sigma_{-i}^0$ ). By continually updating the opponent distribution during play, the adjusted Q-Mixing result better responds to the actual opponent.

An interpretation of Q-Mixing with an OPC is as a decomposition of the problem of learning a best-response to a mixed-strategy opponent. Instead of directly attempting to solve this problem, the problem is broken down into the sub-problems of opponent policy identification and response. Focusing response learning on only an individual opponent policy removes specious experiences caused by misattributing its origin to the wrong opponent policy. An easier reinforcement learning problem is thereby established that reduces the required simulation cost.

An ancillary benefit of the opponent classifier is that poorly estimated Q-values tend to have their impact minimized. For example, if an observation occurs only against the second pure strategy, then the Q-value against the first pure strategy would not be trained well, and thus could distort the policy from Q-Mixing. These poorly trained cases correspond to unlikely opponents and

**Input:**  $\mathcal{O}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \epsilon, \gamma$   
 $V_0(o | \sigma_{-i}) \leftarrow \sum_{\pi_{-i}} \sigma_{-i}(\pi_{-i}) Q(o, a | \pi_{-i})$   
**do**  
 $Q_t(o, a | \sigma_{-i}) \leftarrow \sum_{\pi_{-i}} \psi(\pi_{-i} | o, \sigma_{-i}) \sum_{o', r} \mathcal{T}(o', r | o, a, \pi_{-i}) [r + \gamma V_{t-1}(o' | \sigma_{-i})]$   
 $V_t(o | \sigma_{-i}) \leftarrow \max_a Q_t(o, a | \sigma_{-i})$   
 $\pi_t(o | \sigma_{-i}) \leftarrow \arg \max_a Q_t(o, a | \sigma_{-i})$   
**while**  $\exists o \in \mathcal{O} | V_t(o) - V_{t-1}(o) | > \epsilon$   
**Output:**  $V_t, Q_t, \pi_t$

Algorithm 1. Value iteration: Q-Mixing.

get reduced weighting in the version of Q-Mixing augmented by the classifier.

### 3.3. Accounting for future uncertainty

To account for future uncertainty, Q-Mixing must be able to update its successor observation values given future observations. This can be done by expanding the Q-value into its components: expected reward under the current belief in the opponent's policy, and our expected next observation value. By updating the second term to recursively reference an updated opponent belief we can account for changing beliefs in the future. The extended formulation, Q-Mixing-Value-Iteration (QMVI), is given by:

$$Q_i^*(o_i^t, a_i^t | \sigma_{-i}) = \sum_{\pi_{-i} \in \Pi_{-i}} \psi_i(\pi_{-i} | o_i^t, \sigma_{-i}) \cdot r_i(o_i^t, a_i^t | \pi_{-i}) + \gamma \mathbb{E}_{o_i^{t+1}} [V^*(o_i^{t+1} | \sigma_{-i})].$$

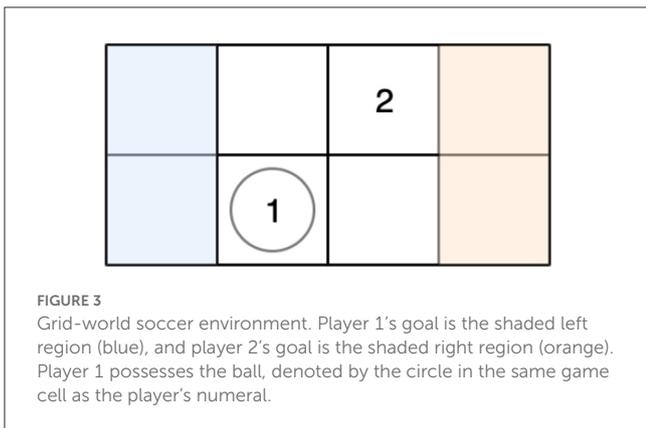
If we assume that we have access to both a dynamics model and the observation distribution dependent on the opponent, then we can directly solve for this quantity through Value Iteration (Algorithm 1). These requirements are quite strong, essentially requiring perfect knowledge of the system with regards to all opponents. The additional step of Value Iteration also carries a computational burden, as it requires iterating over the full observation and action spaces. Though these costs may render QMVI infeasible in practice, we provide Algorithm 1 as a way to ensure correctness in Q-values.

## 4. Experiments

### 4.1. Grid-world soccer

We first evaluate Q-Mixing on a simple grid-world soccer environment (Littman, 1994; Greenwald and Hall, 2003). This environment has small state and action spaces, allowing for inexpensive simulation. With this environment we pose the following questions:

1. Can QMVI obtain Q-values for mixed-strategy opponents?
2. Can Q-Mixing transfer Q-values across all of the opponent's mixed strategies?



The soccer environment is composed of a soccer field, two players, one ball, and two goals. The player's objective is to acquire the ball and score a goal while preventing the opponent from scoring on their own goal. The scorer receives +1 reward, and the defender receives -1 reward. The state of the environment consists of the entire field including the locations of the players and ball. This is represented as a  $5 \times 4$  matrix with six possible values in each cell, referring to the occupancy of the cell. The players may move in any of the four cardinal directions or stay in place. Actions taken by the players are executed in a random order, and if the player possessing the ball moves last then the first player may take possession of the ball by colliding with them. An illustration of the soccer environment can be seen in Figure 3.

In our experiments, we learn policies for Player 1 using Double DQN (van Hasselt et al., 2016). The state space is raveled into a vector of length 120 that is then fed into a deep neural network. The network has two fully-connected hidden layers of size 50 with ReLU activation functions and an output layer over the actions with size 5. Player 2 plays a strategy over five policies, each using the same shape neural networks as Player 1, generated using the double oracle (DO) algorithm (McMahan et al., 2003). These policies are frozen for the duration of the experiments. Further details of the environment and experiments are in the Appendix.

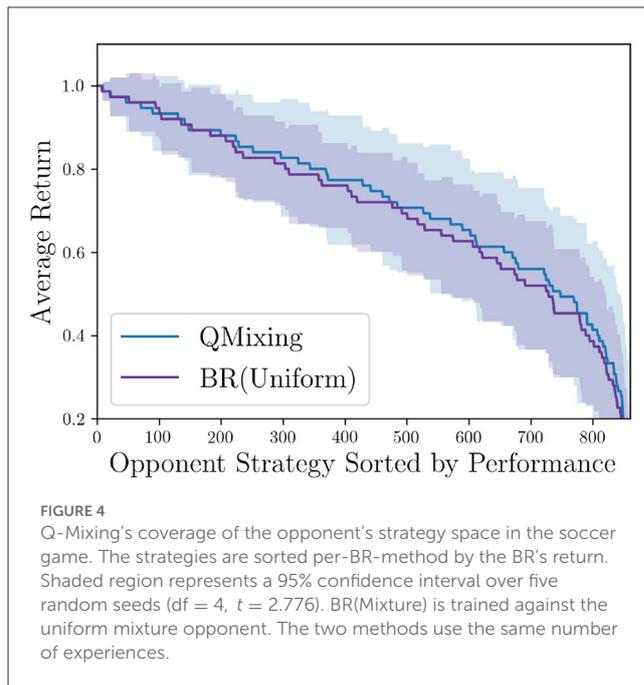
#### 4.1.1. Empirical verification of Q-Mixing

We now turn to our first question: whether QMVI can obtain Q-values for mixed-strategy opponents. To answer this, we run the QMVI algorithm against a fixed opponent mixed strategy (Algorithm 1). We construct dynamics models for each opponent by considering the opponent's policy and the multiagent dynamics model as a single entity. Then we may approximate the relative observation-occupancy distribution by rolling out 30 episodes against each policy and estimating the distribution.

In our experiment, an optimal policy was reached in 14 iterations. The resulting policy best-responded to each individual opponent and the mixture. This empirically validates our first hypothesis.

#### 4.1.2. Coverage of opponent strategy space

Our second question is whether Q-Mixing can produce high-quality responses for any opponent mixture. Our evaluation of



this question employs the same five opponent pure strategies as the previous experiment. We first trained a baseline,  $BR(Uniform)$  or  $BR(\sigma_{-i})$ , directly against the uniform mixed-strategy opponent. DQN's hyperparameters were tuned against a held-out set of evaluation opponent policies. The same hyperparameters were then used to train against each of the opponent's pure strategies, with the simulation budget split equally. The Q-values trained respectively are used as the components for Q-Mixing, and an OPC is also trained from their replay buffers. The OPC has the same neural network architecture as the policies, but modifies the last layer to predict over the size of the opponent's strategy set.

We evaluate each method against all opponent mixtures truncated to the tenths place [e.g., (0.1, 0.6, 0.1, 0.2, 0.0)], resulting in 860 strategy distributions. This is meant to serve as a representative coverage of the entire opponent's mixed-strategy space. For each one of these mixed-strategies, we simulate the performance of each method against that mixture for thirty episodes. We then collect the average payoff against each opponent mixture and sort the averages in descending order.

Figure 4 shows Q-Mixing's performance across the opponent mixed-strategy space. Learning in this domain is fairly easy, so both methods are expected to win against almost every mixed-strategy opponent. Nevertheless, Q-Mixing generalizes across strategies better, albeit with slightly higher variance. While the improvement of Q-Mixing is incremental, we interpret this first evaluation as validating the promise of Q-Mixing for coverage across mixtures.

### 4.2. Sequential social dilemma game

In this section, we evaluate the performance of Q-Mixing in a more complex setting, in the domain of sequential social dilemmas (SSD). These dilemmas present each agent with difficult strategic decision where they must balance collective-

and self-interest throughout repeated interactions. Like our grid-world soccer example, our SSD *Gathering* game has two players, but unlike the soccer game it is general sum. Most importantly, the game exhibits imperfect observation, that is, the environment is only partially observable and the players have private information. We investigated the following research questions on this environment:

1. Can Q-Mixing work with Q-values based on observations from a complex environment?
2. Can the use of an opponent policy classifier that updates the opponent distribution improve performance?
3. Can policy distillation be used to compress a Q-Mixing based policy, while preserving performance?

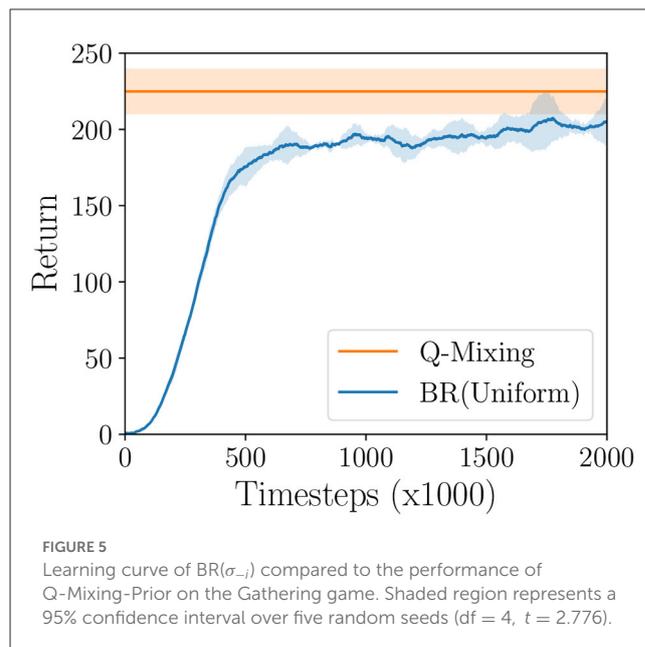
The *Gathering* game is a gridworld instantiation of a tragedy-of-the-commons game (Perolat et al., 2017). In this game, the players compete to harvest apples from an orchard; however, the regrowth rate of the apples is proportional to the number of nearby apples. Limited resources puts the players at odds, because they must work together to maintain the orchard's health while selfishly amassing apples. In addition to picking apples, the players may tag all players in a beam in front of them removing them from the game for a short fixed duration. Moreover, the agents face partial observation as they are only able to see a small window in front of them and do not know the full state of the game.

A visualization of the environment and more details are included in Appendix. The policies used in this environment are implemented as deep neural networks with two hidden layers with 50 units and ReLU activations. Agents can choose to move in the four cardinal directions, turn left or right, tag the other player, or perform no action. The policies are trained using the Double DQN algorithm (van Hasselt et al., 2016), and the opponent's parameters are always kept fixed (no training updates are performed) for the duration of the experiments.

All statistical results below are reported with a 95% confidence interval based on the Student's  $t$ -distribution ( $df = 4$ ,  $t = 2.776$ ). To generate this interval, we perform almost the entire experimental process, as described in each experiment's subsection, entirely across five random seeds. For each seed we do not resample hyperparameters nor generate new opponents. A final result is produced for each random seed, often resulting in a sample of payoffs for a variety of player profiles. Then the sample statistics from each random seed are utilized to construct the confidence interval.

#### 4.2.1. Empirical verification of Q-Mixing

We experimentally evaluate Q-Mixing-Prior on the *Gathering* game, allowing us to confirm our algorithm's robustness to complex environments. First, a set of three opponent policies are generated through DO. A BR is trained against each of the independent pure-strategies. A baseline BR( $\sigma_{-i}$ ), or BR(Uniform), is trained directly against the uniform mixture of the same opponents. We evaluate Q-Mixing-Prior's ability to transfer strategic knowledge by first simulating its performance against the mixed-strategy opponent. Then we compare this performance to the performance that BR( $\sigma_{-i}$ )'s



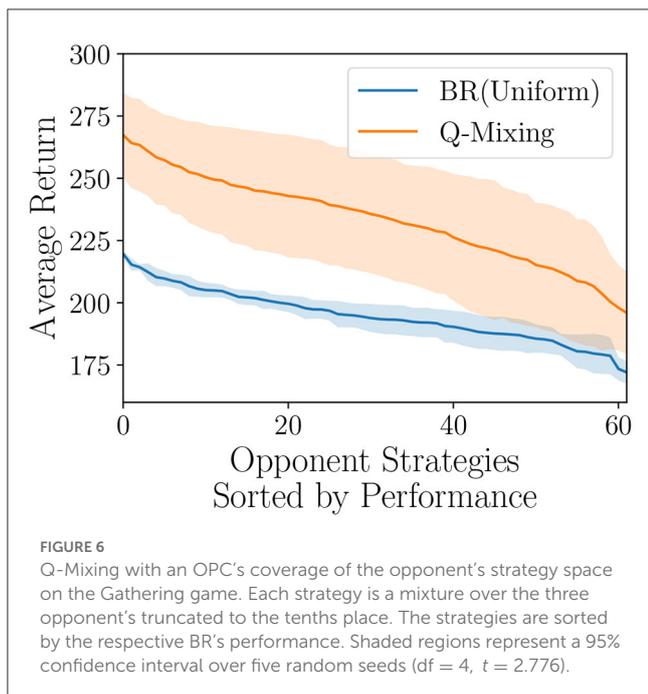
strategy that was learned by training directly against the mixed-strategy.

The training curves for BR( $\sigma_{-i}$ ) is presented in Figure 5. From this graph we can see that Q-Mixing-Prior is able to achieve performance stronger than BR( $\sigma_{-i}$ ) through transfer learning. This is possible because the pure-strategy BR are able to learn stronger policies from specialization compared to the mixed-strategy BR. For example, when playing only against  $\pi_{-i}^0$  the best-response BR( $\pi_{-i}^0$ ) achieves a return of  $227.93 \pm 21.01$  while BR( $\sigma_{-i}$ ) achieves  $207.98 \pm 15.92$ . This verifies our hypothesis that Q-Mixing is able to transfer knowledge under partial observation.

#### 4.2.2. Opponent classification

Our next research question is: can the use of an OPC that updates the opponent distribution in Q-Mixing improve its performance? During play against an opponent sampled from the mixed-strategy, the player is able to gather evidence about which opponent they are playing. We hypothesize that leveraging this evidence to weight the importance of the respective BR's Q-values will improve Q-Mixing's performance.

To verify this hypothesis, we train an OPC using the replay buffers associated with each BR policy. These are the same buffers that were used to train the BRs, and cost no additional compute to collect. The replay buffers are concatenated and labeled with their source policy and constitute our classification dataset. From this data, our classifier must produce a distribution over opponent pure strategies given an observation. The OPC is implemented with a deep neural network two hidden layers with 50 units and ReLU activations. This is the same neural network architecture used by the policies. To train this classifier we take each experience from the pure-strategy BR replay buffers and assign them each a class label for each respective pure-strategy. The classifier is then trained to predict this label using a cross-entropy loss.

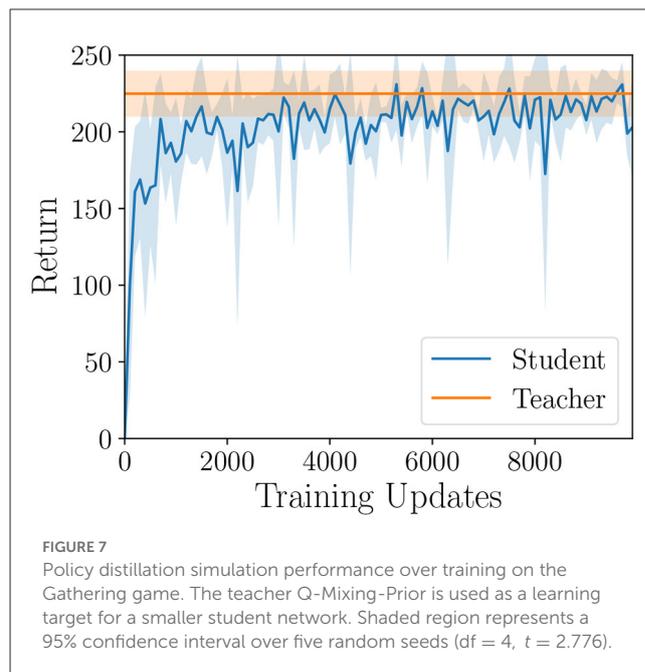


We evaluate Q-Mixing-OPC by testing the performance on a representative coverage of the mixed-strategy opponents illustrated in Figure 6. We found that the Q-Mixing-OPC policy performed stronger against the full opponent strategy coverage supporting our hypothesis that an OPC can identify the opponent's pure-strategy and enable Q-Mixing to choose the correct BR policy. However, our method has a much larger variance, which can be explained by the OPC's misclassification of the opponent resulting in poorly chosen actions throughout the episode. This first result in opponent classification suggests further investigation into stronger classification methods is a fruitful direction for future research.

#### 4.2.3. Policy distillation

In Q-Mixing we need to compute  $Q$ -values for each of the opponent's pure strategies. This can be a limiting factor in parametric policies, like deep neural networks, where our policy's complexity grows linearly in the size of the support of the opponent's mixture. This can become unwieldy in both memory and computation. To remedy these issues, we propose using policy distillation to compress a Q-Mixing policy into a smaller parametric space (Hinton et al., 2014).

In the policy distillation framework, a larger neural network referred to as the *teacher* is used as a training target for a smaller neural network called the *student*. In our experiment, the Q-Mixing policy is the teacher to a student neural network that is the size of a single best-response policy. The student is trained in a supervised learning framework, where the dataset is the concatenated replay buffers from training pure-strategy best-responses. This is the same dataset that was used in opponent classifying, which was notably generated without running any additional simulations. A batch of data is sampled from the replay-buffer and the student predicts  $Q^S$  the teacher's  $Q^T$   $Q$ -values for each action. The student is then



trained to minimize the KL-divergence between the predicted  $Q$ -values and the teacher's true  $Q$ -values. There is a small wrinkle, the policies produce  $Q$ -values, and KL-divergence is a metric over probability distributions. To make this loss function compatible, the  $Q$ -values are transformed into a probability distribution by softmax with temperature  $\tau$ . The temperature parameter allows us to control the softness of the maximum operator. A high temperature produces actions that have a near-uniform probability, and as the temperature is lowered the distribution concentrates weight on the highest  $Q$ -Values (Sutton and Barto, 2018). The benefit of a higher temperature is that more information can be passed from the teacher to the student about each state. Additional training details are described in the Supplementary material.

The learning curve of the student is reported in Figure 7. We found that the student policy was able to recover the performance of Q-Mixing-Prior, albeit with slightly higher variance. This study did not include any effort to optimize the student's performance, thus further improvements with the same methodology may be possible. This result confirms our hypothesis that policy distillation is able to effectively compress a policy derived by Q-Mixing.

## 5. Related work

### 5.1. Multiagent learning

The most relevant line of work in multiagent learning studies the interplay between centralized learning and decentralized execution (Tan, 1993; Kraemer and Banerjee, 2016). In this regime, agents are able to train with additional information about the other agents that would not be available during evaluation (e.g., the other agent's state Rashid et al., 2018, actions Claus and Boutilier, 1998, or a coordination signal Greenwald and Hall, 2003). The key question then becomes: how to create a policy that can be evaluated without additional information? A popular approach is

to decompose the joint-action value into independent Q-values for each agent (Guestin et al., 2001; He et al., 2016; Rashid et al., 2018; Sunehag et al., 2018; Mahajan et al., 2019). An alternative approach is to learn a centralized critic, which can train independent agent policies (Gupta et al., 2017; Lowe et al., 2017; Foerster J. N. et al., 2018). Some work has proposed constructing meta-data about the agent's current policies as a way to reduce the learning instability present in environments where other agents' policies are changing (Foerster et al., 2017; Omidshafiei et al., 2017).

A set of assumptions that can be made is that all players have fixed strategy sets. Under these assumptions, agents could maintain more sophisticated beliefs about their opponent (Zheng et al., 2018) and extend this to recursive-reasoning procedures (Yang et al., 2019). These lines of work focus more on other-player policy identification and are a promising future direction for improving the quality of the OPC. One more potential extension of the OPC is to consider alternative objectives. Instead of focusing exclusively on predicting the opponent, in safety-critical situations an agent will want to consider an objective that accounts for inaccurate prediction of their opponent. The Restricted Nash Response Johanson et al. (2007) encapsulates this measure by balancing maximal performance if the prediction is correct balanced with reasonable performance if the prediction is inaccurate. While both of these directions of work focus around opponent-policy prediction, they do so under a largely different problem statement. Most notably, these works do not consider varying the distribution of the opponent policies, and as such extending these works to fit this new problem statement would constitute its own study.

Instead of building or using complete models of opponents, one may use an implicit representation of their opponents. By choosing to build an explicit model of their opponent they circumvent needing a large amount of data to reconstruct the opponent policy. An additional benefit is that there are less likely to be errors in the model that need to be overcome, because a perfect reconstruction of a complex policy is no longer necessary. He et al. (2016) proposes DRON which uses a learned latent action prediction of the opponent as conditioning information to the policy (in a similar nature to the opponent-actions in the joint-action value area). They also show another version DRON which uses a Mixture-of-Experts (Jacobs et al., 1991) operation to marginalize over the possible opponent behaviors. More formally, they compute the marginal  $\sum_{a_{-i}} \pi_{-i}(a_{-i}|s_i) Q_i(s_i, a_i, a_{-i})$ , which is over the action-space and utilized to condition the expected Q-value. Q-Mixing is built off a similar style of marginalization; however, it marginalizes over the policy-space of the opponent instead of the action-space. Moreover, Q-Mixing depends on independent BR Q-values against each opponent-policy, where DRON learns a single Q-network. Bard et al. (2013) proposes implicitly modeling opponents through the payoffs received from playing against a portfolio of the agent's policies.

Most multiagent learning work focuses on the simultaneous learning of many agents, where there is not a distribution over a static set of opponent policies. This difference in methods can have strong influences on the final learned policies. For example, when a policy is trained concurrently with another particular opponent-policy they may overfit to each other's behavior. As a result, the final learned policy may be unable to coordinate or play well against

any other opponent policy (Bard et al., 2020). Another potential problem is that each agent now faces a dynamic learning problem, where they must learn a moving target (the other agent's policy; Tesauro, 2003; Foerster J. et al., 2018).

## 5.2. Multi-task learning

Multiagent learning is analogous to multi-task learning. In this reconstruction, each strategy/policy is analogous to solving a different task. And the opponent's strategy would be the distribution over tasks. Similar analogies to tasks can be made with objectives, goals, contexts, etc. (Kaelbling, 1993; Ruder, 2017).

The multi-task community has roughly separated learnable knowledge into two categories (Snel and Whiteson, 2014). *Task relevant* knowledge pertains to a particular task (Jong and Stone, 2005; Walsh et al., 2006); meanwhile, *domain relevant* knowledge is common across all tasks (Caruana, 1997; Foster and Dayan, 2002; Konidaris and Barto, 2006). Work has been done that bridges the gap between these settings; for example, knowledge about a task could be a curriculum to utilize over tasks (Czarnecki et al., 2018). In task relevant learning, a leading method is to identify state information that is irrelevant to decision making, and abstract it away (Jong and Stone, 2005; Walsh et al., 2006). Our work falls into the same task relevant category, where we are interested in learning responses to particular opponent policies. What differentiates our work from the previous work is that we learn Q-values for each task independently, and do not ignore any information.

Progressively growing neural networks is another similar line of work (Rusu et al., 2016), focused on a stream of new tasks. Schwarz et al. (2018) also found that network growth could be handled with policy distillation.

## 5.3. Transfer learning

Transfer learning is the study of reusing knowledge to learn new tasks/domains/policies. Within transfer learning, we look at either how knowledge is transferred, or what kind of knowledge is transferred. Previous work on how to transfer knowledge has tended to follow one of two main directions (Pan and Yang, 2010; Lampinen and Ganguli, 2019). The *representation transfer* direction considers how to abstract away general characteristics about the task that are likely to apply to later problems. Ammar et al. (2015) present an algorithm where an agent collects a shared general set of knowledge that can be used for each particular task. The second direction directly transfers parameters across tasks; appropriately called *parameter transfer*. Taylor et al. (2005) show how policies can be reused by creating a projection across different tasks' state and action spaces.

In the literature, transferring knowledge about the opponent's strategy is considered intra-agent transfer (Silva and Costa, 2019). The focus of this area is on *adapting to other agents*. One line of work in this area focuses on *ad hoc* teamwork, where an agent must learn to quickly interact with new teammates (Barrett and Stone, 2015; Bard et al., 2020). The main approach relies on already having a set of policies available, and learning to select which

policy will work best with the new team (Barrett and Stone, 2015). Another work proposes learning features that are independent of the game, which can either be qualities general to all games or strategies (Banerjee and Stone, 2007). Our study differs from these in its focus on the opponent's policies as the source of information to transfer.

## 6. Conclusions

This paper introduces Q-Mixing, an algorithm for transferring knowledge across distributions of opponents. We show how Q-Mixing relies on the theoretical relationship between an agent's action-values, and the strategy employed by the other agents. A first empirical confirmation of the approach is demonstrated using a simple grid-world soccer environment. In this environment, we show how experience against pure strategies can transfer to construction of policies against mixed-strategy opponents. Moreover, we show that this transfer is able to cover the space of mixed strategies with no additional computation.

Next, we tested our algorithm's robustness on a sequential social dilemma. In this environment, we show the benefit of introducing an opponent policy classifier, which uses the agent's observations to update its belief about the opponent's policy. This updated belief is then used to revise the weighting of the respective SRQVs.

Finally, we address the concern that a Q-Mixing policy may become too large or computationally expensive to use. To ease this concern we demonstrate that policy distillation can be used to compress a Q-Mixing policy into a much smaller parameter space.

## Data availability statement

The original contributions presented in the study are included in the article/Supplementary material, further inquiries can be directed to the corresponding author.

## References

- Ammar, H. B., Eaton, E., Luna, J. M., and Ruvolo, P. (2015). "Autonomous cross-domain knowledge transfer in lifelong policy gradient reinforcement learning," in *24th International Conference on Artificial Intelligence*. (Buenos Aires: IJCAI), 3345–3349.
- Banerjee, B., and Stone, P. (2007). "General game learning using knowledge transfer," in *20th International Joint Conference on Artificial Intelligence* (Hyderabad: IJCAI), 672–677.
- Bard, N., Foerster, J. N., Chandar, S., Burch, N., Lanctot, M., Song, H. F., et al. (2020). The Hanabi challenge: a new frontier for AI research. *Artif. Intell.* 280, 506. doi: 10.48550/arXiv.1902.00506
- Bard, N., Johanson, M., Burch, N., and Bowling, M. (2013). "Online implicit agent modelling," in *12th International Conference on Autonomous Agents and Multiagent Systems* (Saint Paul, MI), 255–262.
- Barrett, S., and Stone, P. (2015). "Cooperating with unknown teammates in complex domains: a robot soccer case study of *ad hoc* teamwork," in *29th AAAI Conference on Artificial Intelligence*. (Austin: AAAI), 2010–2016.
- Caruana, R. (1997). Multitask learning. *Machine Learn.* 28, 41–75.
- Claus, C., and Boutillier, C. (1998). "The dynamics of reinforcement learning in cooperative multiagent systems," in *15th National Conference on Artificial Intelligence*. (Madison: AAAI), 746–752.
- Czarnecki, W., Jayakumar, S., Jaderberg, M., Hasenclever, L., Teh, Y. W., Heess, N., et al. (2018). "Mix & match agent curricula for reinforcement learning," in *35th International Conference on Machine Learning*. (Stockholm: ICML), 1087–1095.
- Foerster, J., Chen, R. Y., Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I. (2018). "Learning with opponent-learning awareness," in *17th International Conference on Autonomous Agents and Multiagent Systems*. (Stockholm: AAMAS), 122–130.
- Foerster, J., Nardelli, N., Farquhar, G., Afouras, T., Torr, P. H. S., Kohli, P., et al. (2017). "Stabilising experience replay for deep multi-agent reinforcement learning," in *34th International Conference on Machine Learning*. (Sydney: ICML), 1146–1155.
- Foerster, J. N., Farquhar, G., Afouras, T., Nardelli, N., and Whiteson, S. (2018). "Counterfactual multi-agent policy gradients," in *32nd AAAI Conference on Artificial Intelligence*. (New Orleans, LA: AAAI), 2974–2982.
- Foster, D., and Dayan, P. (2002). Structure in the space of value functions. *Machine Learn.* 49, 325–346. doi: 10.1023/A:1017944732463

## Author contributions

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

## Funding

Research for this article conducted at the University of Michigan was supported in part by funding from the US Army Research Office (MURI grant: W911NF-13-1-0421) and from DARPA SI3-CMD.

## Conflict of interest

TA was employed by the company DeepMind Technologies Ltd.

The remaining 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.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2023.804682/full#supplementary-material>

- Greenwald, A., and Hall, K. (2003). "Correlated-Q learning," in *20th International Conference on Machine Learning*. (Washington, DC: ICML), 242–249.
- Guestin, C., Koller, D., and Parr, R. (2001). "Multiagent planning with factored MDPs," in *14th International Conference on Neural Information Processing Systems* (Vancouver: NeurIPS), 1523–1530.
- Gupta, J. K., Egorov, M., and Kochenderfer, M. (2017). "Cooperative multi-agent control using deep reinforcement learning," in *16th International Conference on Autonomous Agents and Multiagent Systems*, eds G. Sukthankar and J. A. Rodriguez-Aguilar (São Paulo: AAMAS), 66–83.
- He, H., Boyd-Graber, J., Kwok, K., and III, H. D. (2016). "Opponent modeling in deep reinforcement learning," in *33rd International Conference on Machine Learning*. (New York City, NY: ICML), 1804–1813.
- Hinton, G., Vinyals, O., and Dean, J. (2014). "Distilling the knowledge in a neural network," in *NeurIPS: Deep Learning and Representation Learning Workshop* (Montréal).
- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Comput.* 3, 79–87.
- Jaderberg, M., Czarnecki, W. M., Dunning, I., Marris, L., Lever, G., Castañeda, A. G., et al. (2019). Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science* 364, 859–865. doi: 10.1126/science.aau6249
- Johanson, M., Zinkevich, M., and Bowling, M. (2007). "Computing robust counter-strategies," in *Advances in Neural Information Processing Systems 20*, eds J. Platt, D. Koller, Y. Singer and S. Roweis (Vancouver: Curran Associates, Inc.).
- Jong, N. K., and Stone, P. (2005). "State abstraction discovery from irrelevant state variables," in *19th International Joint Conference on Artificial Intelligence*. (Edinburgh: IJCAI), 752–757.
- Kaelbling, L. P. (1993). "Learning to achieve goals," in *13th International Joint Conference on Artificial Intelligence*. (Chambéry: IJCAI), 1094–1099.
- Konidaris, G., and Barto, A. (2006). "Autonomous shaping: Knowledge transfer in reinforcement learning," in *23rd International Conference on Machine Learning*. (Pittsburgh: ICML), 489–496.
- Kraemer, L., and Banerjee, B. (2016). Multi-agent reinforcement learning as a rehearsal for decentralized planning. *Neurocomputing* 190, 82–94. doi: 10.1016/j.neucom.2016.01.031
- Lampinen, A. K., and Ganguli, S. (2019). "An analytic theory of generalization dynamics and transfer learning in deep linear networks," in *7th International Conference on Learning Representations*. New Orleans: ICLR.
- Lancot, M., Zambaldi, V., Gruslys, A., Lazaridou, A., Tuyls, K., Pérolat, J., et al. (2017). "A unified game-theoretic approach to multiagent reinforcement learning," in *31st International Conference on Neural Information Processing Systems*. (Long Beach, CA: NuerIPS), 4193–4206.
- Littman, M. L. (1994). "Markov games as a framework for multi-agent reinforcement learning," in *11th International Conference on International Conference on Machine Learning*. (New Brunswick, NJ: ICML), 157–163.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., and Mordatch, I. (2017). "Multi-agent actor-critic for mixed cooperative-competitive environments," in *31st International Conference on Neural Information Processing Systems*. (Long Beach, CA: NeurIPS), 6382–6393.
- Mahajan, A., Rashid, T., Samvelyan, M., and Whiteson, S. (2019). "Maven: multi-agent variational exploration," in *33rd Conference on Neural Information Processing Systems*. Vancouver: NeurIPS.
- McMahan, H. B., Gordon, G. J., and Blum, A. (2003). "Planning in the presence of cost functions controlled by an adversary," in *20th International Conference on International Conference on Machine Learning*. (Washington, DC: ICML), 536–543.
- Omidshafiei, S., Papis, J., Amato, C., How, J. P., and Vian, J. (2017). "Deep decentralized multi-task multi-agent reinforcement learning under partial observability," in *34th International Conference on Machine Learning*. Sydney: ICML.
- Pan, S. J., and Yang, Q. (2010). A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* 22, 1345–1359. doi: 10.1109/TKDE.2009.191
- Pérolat, J., Z. Leibo, J., Zambaldi, V., Beattie, C., Tuyls, K., and Graepel, T. (2017). "A multi-agent reinforcement learning model of common-pool resource appropriation," in *31st Conference on Neural Information Processing Systems* (Long Beach, CA).
- Rashid, T., Samvelyan, M., Witt, C. S. d., Farquhar, G., Foerster, J., and Whiteson, S. (2018). "QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning," in *35th International Conference on Machine Learning*. (Stockholm: ICML), 4295–4304.
- Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *CoRR* 2017, abs/1706.05098. doi: 10.48550/arXiv.1706.05098
- Rusu, A. A., Colmenarejo, S. G., Gulcehre, C., Desjardins, G., Kirkpatrick, J., Pascanu, R., et al. (2015). "Policy distillation," in *International Conference on Learning Representations*. (San Diego, CA: ICLR).
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., et al. (2016). Progressive neural networks. *CoRR* 2016, abs/1606.04671. doi: 10.48550/arXiv.1606.04671
- Schwarz, J., Luketina, J., Czarnecki, W. M., Grabska-Barwinska, A., Teh, Y. W., Pascanu, R., et al. (2018). "Progress & compress: a scalable framework for continual learning," in *35th International Conference on Machine Learning*. Stockholm: ICML.
- Silva, F., and Costa, A. (2019). A survey on transfer learning for multiagent reinforcement learning systems. *J. Artif. Intell. Res.* 64, 11396. doi: 10.1613/jair.1.11396
- Smith, M. O., Anthony, T., and Wellman, M. P. (2021). "Iterative empirical game solving via single policy best response," in *9th International Conference on Learning Representations* (Virtual).
- Snel, M., and Whiteson, S. (2014). "Learning potential functions and their representations for multi-task reinforcement learning," in *13th International Conference on Autonomous Agents and Multiagent Systems*. (Paris: AAMAS), 637–681.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., et al. (2018). "Value-decomposition networks for cooperative multi-agent learning," in *17th International Conference on Autonomous Agents and Multiagent Systems*. Stockholm: AAMAS.
- Sutton, R. S., and Barto, A. G. (2018). *Reinforcement Learning: An Introduction, 2nd Edn*. Cambridge, MA: MIT Press.
- Tan, M. (1993). "Multi-agent reinforcement learning: Independent vs. cooperative agents," in *10th International Conference on Machine Learning*. (Amherst, MA: ICML), 330–337.
- Taylor, M. E., Stone, P., and Liu, Y. (2005). "Value functions for RL-based behavior transfer: a comparative study," in *20th National Conference on Artificial Intelligence*. (Palo Alto: AAAI), 880–885.
- Tesauro, G. (2003). "Extending Q-learning to general adaptive multi-agent systems," in *16th International Conference on Neural Information Processing Systems*. (Vancouver: NeurIPS), 871–878.
- van Hasselt, H., Guez, A., and Silver, D. (2016). "Deep reinforcement learning with double Q-learning," in *30th AAAI Conference on Artificial Intelligence*. (Phoenix: AAAI), 2094–2100.
- Walsh, T. J., Li, L., and Littman, M. L. (2006). "Transferring state abstractions between MDPs," in *ICML-06 Workshop on Structural Knowledge Transfer for Machine Learning* (Pittsburgh).
- Wang, X., and Sandholm, T. (2003). "Learning near-Pareto-optimal conventions in polynomial time," in *16th International Conference on Neural Information Processing Systems*. (Vancouver: NeurIPS), 863–870.
- Yang, T., Hao, J., Meng, Z., Zhang, C., Zheng, Y., and Zheng, Z. (2019). "Efficient detection and optimal response against sophisticated opponents," in *International Joint Conference on Artificial Intelligence* (Macao).
- Zheng, Y., Meng, Z., Hao, J., Zhang, Z., Yang, T., and Fan, C. (2018). "A deep bayesian policy reuse approach against non-stationary agents," in *32nd Conference on Neural Information Processing Systems* (Montréal).