Selectively Sharing Experiences Improves Multi-Agent Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

We present a novel multi-agent RL approach, *Selective Multi-Agent PER*, in which agents share with other agents a limited number of transitions they observe during training. They follow a similar heuristic as is used in (single-agent) Prioritized Experience Replay, and choose those transitions based on their td-error. The intuition behind this is that even a small number of relevant experiences from other agents could help each agent learn. Unlike many other multi-agent RL algorithms, this approach allows for largely decentralized training, requiring only a limited communication channel between agents. We show that our approach outperforms baseline no-sharing decentralized training. Further, sharing only a small number of experiences outperforms sharing all experiences between agents, and the performance uplift from selective experience sharing is robust across a range of hyperparameters.

1 Introduction

Multi-Agent Reinforcement Learning (RL) is often considered a hard problem: The environment dynamics and returns depend on the joint actions of all agents, leading to significant variance and non-stationarity in the experiences of each individual agent. Much recent work [21, 15] in multi-agent RL has focused on mitigating the impact of these. Our work goes into a slightly different direction: Leveraging the presence of other agents to collaboratively explore the environment more quickly.

We present a novel multi-agent RL approach that allows agents to share a small number of experiences with other agents. The intuition is that if one agent discovers something important in the environment, then sharing this with the other agents should help them learn faster. However, it is crucial that only important experiences are shared - we will see that sharing all experiences indiscriminately will not improve learning. To this end, we take inspiration from a well-established technique in single-agent RL, *prioritized experience replay* (PER). With PER an off-policy algorithm such as DQN will sample experiences not uniformly, but proportionally to "how far off" the current policy's predictions are in each state, formally the *temporal difference (td) error* . We use this same metric to prioritize which experiences to share with other agents.

We dub the resulting multiagent RL approach "selective multiagent PER" or "suPER". In this, agents independently use an off-policy RL algorithm to learn, but with a twist: Each agent relays its highest-td-error experiences to the other agents, who insert them directly into their replay buffer, which they use for learning. This approach has several advantages:

1. It consistently leads to faster learning and higher eventual performance, across hyperparameter settings.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

2. Unlike many "centralized training, decentralized execution" approaches, the suPER learning paradigm allows for (semi-) decentralized training, requiring only a limited bandwidth communication channel between agents.
3. The paradigm is agnostic to the underlying decentralized training algorithm, and can enhance any existing DQN variation including modern algorithms such as Rainbow [8] and other off-policy algorithms.

In the remainder of the paper, we will discuss related literature and technical preliminaries; introduce our novel algorithm in detail; describe experimental evaluation and results; and suggest avenues for future work.

2 Related Work

Multi-agent RL approaches in the literature can be broadly categorized according to the degree of awareness of each learning agent of the other agents in the system [28]. On one end of this spectrum is independent or decentralized learning where multiple learning agents (in the same environment) are unaware or agnostic to the presence of other agents in the system [27, 12]. From the perspective of each agent, learning regards the other agents as part of the non-stationary environment. At the opposing end of the spectrum, in centralized control a single policy controls all agents. In between these two, a number of related strands of research have emerged.

Approaches nearer the decentralized end of the spectrum often use communication between the agents [35, 3, 7]. Several forms of cooperative communication have been formulated, by which agents can communicate various types of messages, either to all agents or to specific agent groups through dedicated channels [11, 24, 18, 19]. While communication among agents could help with coordination, training emergent communication protocols also remains a challenging problem; recent empirical results underscore the difficulty of learning meaningful emergent communication protocols, even when relying on centralized training [10]. Related to this, and often used in conjunction, is modeling other agents [10, 4, 33], which equips agent with some model of other agent's behavior.

This in turn is related to a number of recent approaches using centralized training but decentralized execution. A prevailing paradigm within this line of work assumes a training stage during which a shared network (such as a critic) can be accessed by all agents to learn decentralised (locally-executable) agent policies [15, 1, 20, 4, 21]. These approaches successfully reduce variance during training, e.g. through a shared critic accounting for other agents' behavior, but rely on joint observations and actions of all agents. Within this paradigm, and perhaps closest related to our own work, [2] introduces an approach in which agents share (all of their) experiences with other agents. While this is based on an on-policy actor-critic algorithm, it uses importance sampling to incorporate the off-policy data from other agents, and the authors show that this can lead to improved performance in sparse-reward settings. Our work also relies on sharing experiences among agents, but crucially relies on selectively sharing only some experiences.

The approach we present in this paper relies intrinsically on off-policy RL algorithms. Most notable in this class is DQN [16], which achieved human-level performance on a wide variety of Atari 2600 games. Various improvements have been made to this algorithm since then, including dueling DQN [31], Double DQN (DDQN) [30], Rainbow [8] and Ape-X [9]. Prioritized Experience Replay [22] improves the performance of many of these, and is closely related to our own work. Off-policy actor-critic algorithms [6, 14] have also been developed, in part to extend the paradigm to continuous control domains.

3 Preliminaries

Reinforcement learning (RL) deals with learning optimal policies for sequential decision making in environments for which the dynamics are not fully known [25]. *Multi-Agent Reinforcement Learning* extends RL to multi-agent settings. A common model is a *Markov game*, or *stochastic game* defined as a tuple $\langle S, \mathcal{A}, \mathcal{R}, \mathcal{T}, \gamma \rangle$ with states S, *joint actions* $\mathcal{A} = \{A^i\}_{i=1}^n$ as a collection of action sets A^i, one for each of the n agents, $\mathcal{R} = \{R^i\}_{i=1}^n$ as a collection of reward functions R^i defining the reward $r^i(a_t, s_t)$ that each agent receives when the joint action $a_t \in \mathcal{A}$ is performed at state s_t, and $\mathcal{T}$ as the probability distribution over next states when a joint action is performed.

In the partially observable case, the definition also includes a joint observation function, defining the observation for each agent at each state. In this framework, at each time step an agent has an *experience* $e =< S_t, A_t, R_{t+1} S_{t+1} >$, where the agent performs action A_t at state S_t after which reward R_{t+1} is received and the next state is S_{t+1}. We focus here on decentralized execution settings in which each agent follows its own individual policy π_i and seeks to maximize its discounted accumulated return. In the special case where rewards are shared between all agents a Markov Game is also called a *decentralized POMDP* or dec-POMDP. We do not require this assumption, but we do require that the Markov game is symmetric or anonymous, meaning all agents have the same action and observation spaces and the environment reacts identically to their actions.

Among the variety of RL solution approaches [26, 25], we focus here on *value-based methods* that use state and action estimates to find optimal policies. Such methods typically use a value function $V_\pi(s)$ to represent the expected value of following policy π from state s onward, and a *Q-function* $Q_\pi(s, a)$, to represent for a given policy π the expected rewards for performing action a in a given state s and following π thereafter.

Q-learning [32] is a *temporal difference* (td) method that considers the difference in Q values between two consecutive time steps. Formally, the update rule is

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha[R_{t+1} + \gamma \max_{a'} Q(S_{t+1}, a') - Q(S_t, A_t))]$$

Where, α is the *learning rate* or *step size* setting the amount by which values are updated during training. The learned action-value function, Q, directly approximates q^*, the optimal action-value function. During training, in order to guarantee convergence to an optimal policy, Q-learning typically uses an ϵ-greedy selection approach, according to which the best action is chosen with probability $1 - \epsilon$, and a random action, otherwise.

Given an experience e_t, the *td-error* represents the difference between the Q-value estimate and actual reward gained in the transition and the discounted value estimate of the next best actions. Formally,

$$\mathrm{td}(e_t) = |R + \gamma \max_{a'} Q(S', a') - Q(S, A)| \tag{1}$$

In the vanilla version of Q-learning Q-values are stored in a table, which is impractical in many real-world problem due to large state space size. In deep Q-Learning (DQN) [17], the Q-value table is replaced by a function approximator typically modeled using a neural network such that $Q(s, a, \theta) \approx Q^*(s, a)$, where θ denotes the neural network parameters.

Replay Buffer (RB) and Prioritized RB: In their simplest form, RL algorithms use and discard incoming data after updating the policy. To increase efficiency an experience replay buffer (Lin, 1992) was suggested. This is a memory structure that enables online reinforcement learning algorithms to store and replay past experiences by sampling them from memory. This allows mixing experiences from different time steps in to a single policy update, allowing rare experiences to be used more than once.

Evidence show the replay buffer to stabilize training of the value function for DQN [16, 17] and to reduce the amount of experiences required for an RL-agent to complete the learning process and achieve convergence [22].

Initial approaches that used a replay buffer, uniformly sampled experiences from the buffer. However, some transitions are more effective for the learning process of an RL agents then others [23]. *Prioritized Experience Replay (PER)* [22] explores the idea that replaying and learning from some transitions, rather than others, can enhance the learning process. PER suggests replacing the standard sampling method, where transitions are replayed according to the frequency they were collected from the environment, with a td-error based method, where transitions are sampled according to the value of their td-error.

As a further extension, **stochastic prioritization** balances between strictly greedy prioritization and uniform random sampling. Hence, the probability of sampling transition i is defined as:

$$P(i) = \frac{p_i^\alpha}{\sum_k p_k^\alpha} \tag{2}$$

where $p_i > 0$ is the priority associated to transition i and α determines the weight of the priority ($\alpha = 0$ is uniform sampling). The value of p_i is determined according to the magnitude of the

td-error, such that $p_i = |\delta_i| + \epsilon$, δ_i is the td-error and ϵ is a small positive constant that guarantees that transitions for which the td-error is zero have a non-zero probability of being sampled from the buffer.

4 suPER: Selective Multi-Agent Prioritized Experience Replay

Our approach is rooted in the same intuition as PER: that not all experiences are equally relevant. We use this insight to help agents learn by sharing between them only a (small) number of their most relevant experiences. Our approach builds on standard off-policy algorithms such as DQN, and adds this experience sharing mechanism between collecting experiences and performing gradient updates, in each iteration of the algorithm:

1. (DQN) Collect a rollout of experiences, and insert each agent's experiences into their own replay buffer.
2. (suPER) Each agent shares their most relevant experiences, which are inserted into all the other agents' replay buffers.
3. (DQN) Each agent samples a minibatch of experiences from their own replay buffer, and performs gradient descent on it.

Steps 1 and 3 are standard DQN [18]; we merely add an additional step between collecting experiences and learning on them. As a corollary, this same approach works for any standard variants of steps 1 and 3, such as dueling DQN [31], DDQN [30], Rainbow [8] and other DQN improvements. Algorithm 1 gives a more detailed listing of this algorithm. Beyond DQN and its variants, suPER can

Algorithm 1 suPER algorithm for DQN

```
for each training iteration do
    Collect a batch of experiences b                          ▷ DQN
    for each agent i do                                       ▷ DQN
        Insert b_i into buffer_i                              ▷ DQN
    end for                                                   ▷ DQN
    for each agent i do                                       ▷ suPER
        Select b*_i ⊆ b_i of experiences to share[1]          ▷ suPER
        for each agent j ≠ i do                               ▷ suPER
            Insert b*_i into buffer_j                         ▷ suPER
        end for                                               ▷ suPER
    end for                                                   ▷ suPER
    for each agent i do                                       ▷ DQN
        Sample a train batch b_i from buffer_i                ▷ DQN
        Learn on train batch b_i                              ▷ DQN
    end for                                                   ▷ DQN
end for
```

[1] See section "Experience Selection"

in principle apply to any off-policy RL algorithm. For instance, in Soft Actor Critic (SAC) [6], the resulting algorithm looks exactly like the DQN variant described in Algorithm 1. The only difference is that the "learn on train batch" step is using the actor-critic model and loss from SAC.

4.1 Experience Selection

We describe here three variants of the suPER algorithm, that differ in how they select experiences to be shared.

Deterministic quantile-based experience selection The learning algorithm keeps a list l of the (absolute) td-errors of its last k experiences ($k = 1500$ by default). For a configured bandwidth β and a new experience e_t, the agent shares the experience if its absolute td-error $|\mathrm{td}(e_t)|$ is at least as large as the $k * \beta$-largest absolute td-error in l. In other words, the agent aims to share the top β-quantile of its experiences, where the quantile is calculated over a sliding window of recent experiences.

$$|\mathrm{td}(e_t) \geq \mathrm{quantile}_\beta(\{e_{t'}\}_{t'=t-k}^{t})$$

Deterministic Gaussian experience selection In this, the learning algorithm calculates the mean μ and variance σ^2 of the (absolute) td-errors of the k most recent experiences ($k = 1500$ by default). It then shares an experience e_t if

$$|\mathrm{td}(e_t)| \geq \mu + c \cdot \sigma^2 \tag{3}$$

where c is a constant chosen such that $1 - \mathrm{cdf}_{\mathcal{N}}(c) = \beta$. In other words, we use the c-quantile of a normal distribution with the (sample) mean and variance of most recent experiences. We include and benchmark this variant for two reasons. One, intuitively, we might want to be more sensitive to clusters of outliers, where using a quantile of the actual data might include only part of the cluster, while a Gaussian model might lead to the entire cluster being included. Two, mean and variance could be computed iteratively without keeping a buffer of recent td-errors, and thereby reducing memory requirements. We aim to benchmark if this approximation impacts performance.

Stochastic weighted experience selection Finally, and most closely related to classical single-agent PER, this variant shares each experience with a probability that's proportional to its absolute td-error. In PER, given a train batch size b, we sample b transitions from the replay buffer without replacement, weighted by each transition's td-error. In suPER, we similarly aim to sample a β fraction of experiences, weighted by their td-errors. However, in order to be able to sample transitions online, we calculate for each experience individually a probability that approximates sampling-without-replacement in expectation. Formally, taking $p_i = |\mathrm{td}(e_i)|$, we broadcast experience e_t with probability

$$p = \min\left(1, \beta \cdot \frac{p_i^\alpha}{\sum_k p_k^\alpha}\right)$$

similarly to equation 2, and taking the sum over a sliding window over recent experiences. It is easy to see that if $\beta = 1/\mathrm{batchsize}$, this is equivalent to sampling a single experience, weighted by td-error, from the sliding window. For larger bandwidth β, this approximates sampling multiple experiences without replacement: If none of the $\beta \cdot \frac{p_i^\alpha}{\sum_k p_k^\alpha}$ terms are greater than 1, this is exact. If we have to truncate any of these terms, we slightly undershoot the desired bandwidth.

In our current experiments, we share experiences and update all quantiles, means and variances once per sample batch, for convenience and performance reasons; However, we could do both online, i.e. after every sampled transition, in real-world distributed deployments. We do not expect this to make a significant difference.

5 Experiments and Results

We evaluate the suPER approach on a number of multiagent benchmark domains.

5.1 Algorithm and control benchmarks

Baseline: Dueling DDQN As a baseline, we use a fully independent DDQN algorithm. We use a modern dueling DDQN implementation from RLlib [13], with DQN hyperparameters largely unchanged. We ran this with a single environment thread from which experiences are collected. The algorithm samples a sequence of experiences using joint actions from all the agents' policies, and then inserts each agent's observation-action-reward-observation transitions into that agent's replay buffer. Each agent's policy is periodically trained using a sample from its own replay buffer, sampled using PER. We refer to this baseline as simply "DQN" in the remainder of this section and in figures.

Baseline: Indiscriminate sharing-all-experiences DQN As a second baseline, we compare in some experiments against a DQN variant that shares *all* experiences among agents. In this variant, whenever a batch of experiences is sampled, every agent's trajectory is inserted into all the agents' replay buffers. This is meant to test whether sharing experiences in general is helpful (e.g. by introducing more variety into each agent's replay buffer, or cycling experiences through the buffer faster), or if it is selectivity in sharing that is helpful. Aside from experience sharing, this algorithm is unchanged from the DQN baseline, i.e. we use a dueling DDQN implementation with the same hyperparameters

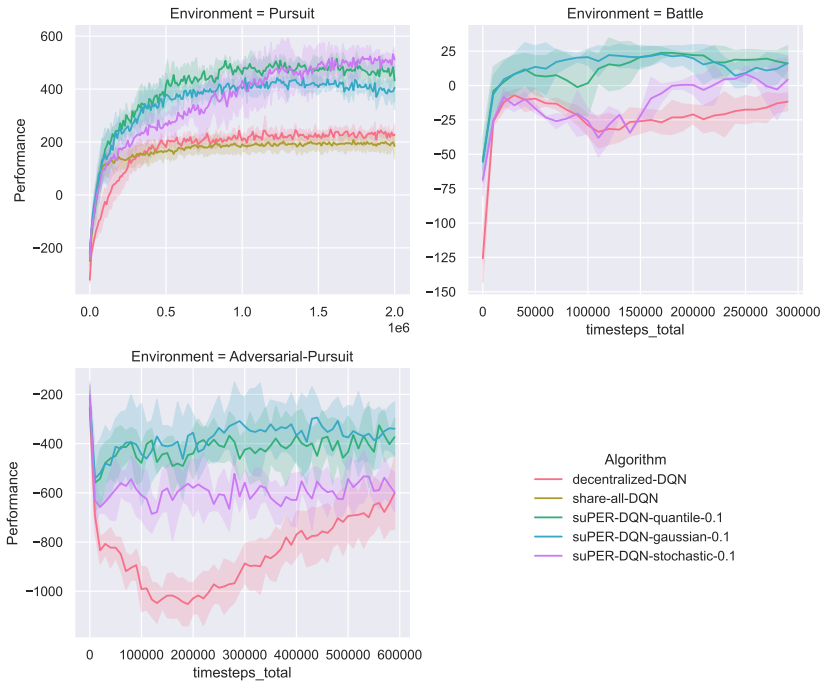

Figure 1: Performance of suPER variants with target bandwidth 0.1 on all three domains. For Pursuit, performance is the total mean episode reward from all agents. For Battle and Adversarial-Pursuit, performance is the total mean episode reward from all agents in the sharing team (blue team in Battle, prey team in Adversarial-Pursuit). Shaded areas indicate one standard deviation.

suPER-DQN We implement suPER on top of the dueling DDQN baseline. Whenever a batch of experiences is sampled, each agent shares only its most relevant experiences (according to one of the criteria described in the previous section) which are then inserted into all other agents' replay buffers. As above, we run the suPER experience sharing on whole sample batches. All DQN hyperparameters are unchanged - the only difference from the baseline is the addition of experience sharing between experience collection and learning. This allows for controlled experiments with like-for-like comparisons.

5.2 Environments

SISL: Pursuit is a semi-cooperative environment, where a group of pursuers has to capture a group of evaders in a grid-world with an obstacle. The evaders (blue) move randomly, while the pursuers (red) are controlled by RL agents. If a group of two or more agents fully surround an evader, they each receive a reward, and the evader is removed from the environment. The episode ends when all evaders have been captured, or after 500 steps, whichever is earlier. Pursuers also receive a (very small) reward for being adjacent to an evader (even if the evader is not fully surrounded), and a (small) negative reward each timestep, to incentivize them to complete episodes early.

MAgent: Battle is a semi-adversarial environment, where two groups of opposing teams are battling against each other. An agent is rewarded 0.2 points for attacking agents in the opposite team, and 5 points if the other agent is killed. All agents start with 10 health points (HP) and lose 2 HP in each attack received, while regaining 0.1 HP in every turn. Once killed, an agent is removed from the environment. An episode ends when all agents from one team are killed. The action space, of size 21 is identical for all agents, with (8) options to attack, (12) to move and one option to do nothing. Since no additional reward is given for collaborating with other agents in the same team, it is considered to be more challenging to form collaboration between agents in this environment.

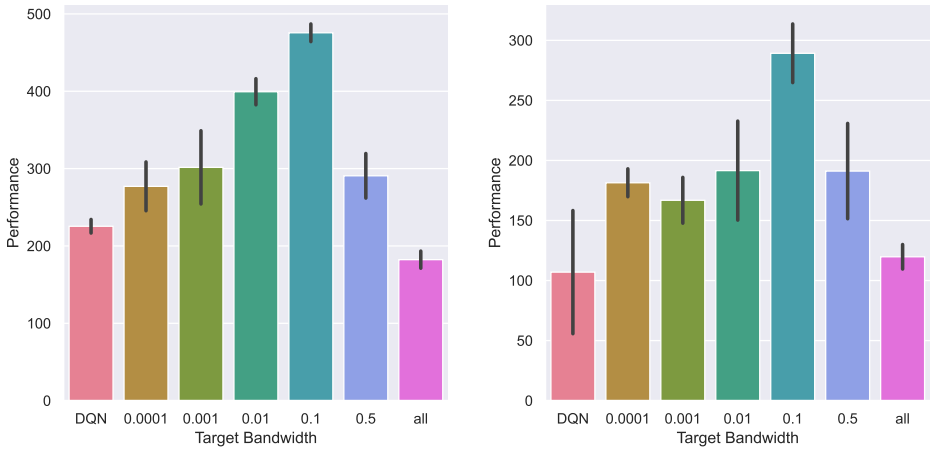

Figure 2: Performance of quantile suPER with varying bandwidth in Pursuit at 1-2M timesteps (left) and at 250k timesteps (right).

MAgent: Adversarial Pursuit is a predator-prey environment, with two types of agents, prey and predator. The predators navigate through obstacles in the map with the purpose of tagging the prey. An agent in the predators team is rewarded 1 point for tagging a prey, while a prey is rewarded -1 when being tagged by a predator. Unlike in the Battle environment, prey agents are not removed from the game when being tagged. Note that prey agents are provided only with a negative or zero reward (when manage to avoid attacks), and their aim is thus to evade predator agents.

5.3 Performance Evaluation

Figure 1 shows learning curves of suPER-DQN, with gaussian and quantile experience selection and target bandwidth 0.1, compared to standard no-sharing DQN. We see that both suPER-DQN variants outperform no-sharing DQN on all three domains. In Pursuit, suPER-DQN achieves over twice the reward of no-sharing DQN at convergence, increasing from 225.9 (std.dev 22.0) to 457.8 (std.dev 55.7) for quantile suPER-DQN measured at 2M training steps. In both Battle and Adversarial-Pursuit, we enabled suPER-sharing for only one of the two teams each, and see that this significantly improved performance of the sharing team (as measured by sum of rewards of all agents in the team acros each episode). In Battle, the blue team performance increase from -11.4 (std.dev 4.8) for no-sharing DQN to 16.2 (std.dev 5.4) for quantile suPER-DQN. In Adversarial-Pursuit, the prey performance increased from -558.7 (std.dev 53.4) to -373.4 (std.dev 77.1). Interestingly, in both of those environment this performance gain did not come purely at the expense of the other, non-sharing team, whose relative performance drop was smaller than the sharing teams' gains in both environments. We hypothesize that suPER may help agents avoid environmental penalties in addition to improving performance against the opposing team. Surprisingly, in the Adversarial-Pursuit domain, sharing between the predator agents had a much smaller effect.

5.4 Bandwidth Sensitivity

In the Pursuit domain, we performed an analysis of the performance of suPER-DQN for varying target bandwidths ranging from 0.0001 to 1 (sharing all experiences). Figure 2 (left) shows the converged performance of suPER-DQN with quantile-based experience selection in Pursuit. A label of "DQN" indicates that no sharing is taking place, i.e. decentralized DQN; a label of "all" indicates that all experiences are shared, without any prioritized selection. Numerical labels give different target bandwidths. Two things stand out to us: First, sharing all experiences indiscriminately does not result in increased performance. In fact, at convergence, it results in slightly lower performance than no-sharing DQN. Second, there is a clear peak of performance around a target bandwidth of 0.01 - 0.1, which also holds for gaussian and stochastic experience selection (we refer the reader to the appendix for more details). We conclude that sharing experiences *selectively* is crucial for learning to benefit from it.

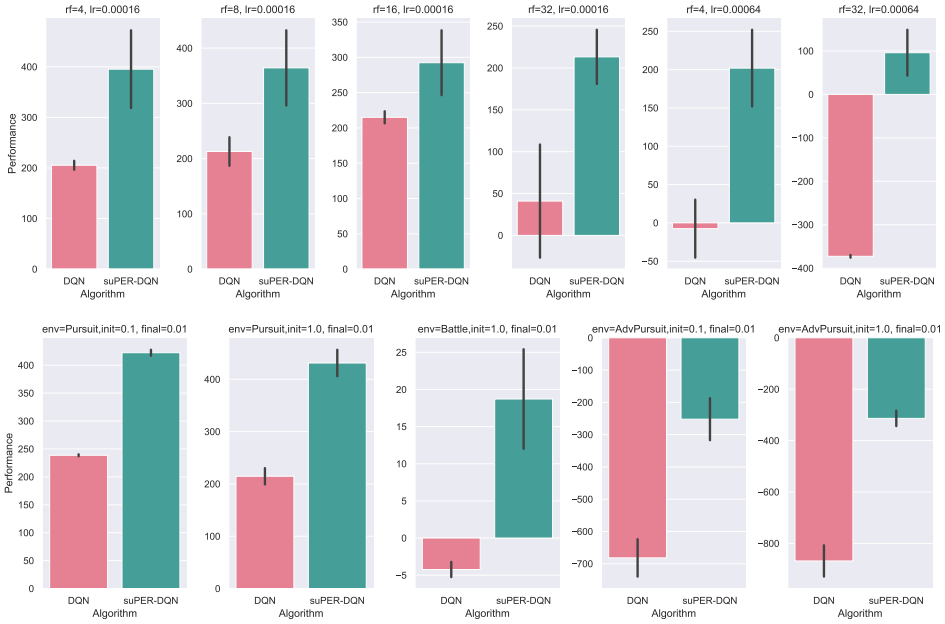

Figure 3: Performance of DQN and suPER-DQN (gaussian experience selection, target bandwidth 0.1) for differing hyperparameter settings of the underlying DQN algorithm. Top: Different learning rates and rollout fragment lengths in Pursuit. Bottom: Different exploration settings in all three environments. Hyperparameters otherwise identical to those used in Figure 1. Performance measured at 1M timesteps in Pursuit, 300k timesteps in Battle, 400k timesteps in Adversarial-Pursuit.

Furthermore, it is noteworthy that early in training, even small bandwidths show significantly higher performance than no-sharing DQN. Figure 2 (right) shows performance at 250k timesteps, where even a target bandwidth of 0.0001 shows significantly improved performance over baseline. This suggests the applicability of suPER to settings where agents need to adapt rapidly to an unknown environment, even if communication is very limited. Indiscriminate sharing of all experiences again did not show a performance benefit at this stage of training. It did show increased performance (similar to selective sharing) very early in training, which we attribute to the faster replay buffer fill rate it entails. After replay buffers had been filled this performance improvement vanished.

5.5 Experience Selection

We find that at bandwidth 0.1, all three experience selection variants performed similarly at convergence. Stochastic experience selection however consistently converge slower than both deterministic variants and had lower performance early in training (but still higher than no-sharing baseline). Gaussian experience selection performed similarly to the quantile selection we designed it to approximate. Its actual used bandwidth was however much less responsive to target bandwidth than the other two variants (e.g. in Pursuit at target bandwith 0.001, Gaussian suPER consumed 0.018 actual bandwidth). We believe this demonstrates that in principle approximating the actual distribution of td-errors using mean and standard deviation is feasible, but that more work is needed in determining the optimal value of c in equation 3.

5.6 Stability across hyperparameters

We evaluate suPER across a range of hyperparameter settings such as differing train batch sizes and exploration configurations. Figure 3 shows performance of no-sharing DQN and suPER-DQN for different hyperparameters. As we can see, suPER-DQN outperforms no-sharing DQN consistently across all the hyperparameter settings considered.

6 Discussion

We present selective multiagent PER, a selective experience-sharing mechanism that can improve off-policy RL algorithm in multiagent settings. Conceptually, our approach is rooted in the same intuition that Prioritized Experience Replay was based on, which is that td-error is a useful approximation of how much an agent could learn from a particular experience. Experimental evaluation on dueling DDQN shows improved performance (as measured in sample efficiency and/or converged performance) across a range of hyperparameters of the underlying algorithm, and in multiple benchmark domains. We generally see significantly higher performance with suPER and a target bandwidth of 0.01-0.1 late in training, suggesting that the approach can help agents learn to coordinate better. Given that this effect appeared consistently across a wide range of hyperparameters and multiple environments, the suPER approach may be useful as a general-purpose multi-agent RL technique. Equally noteworthy is a significantly improved performance early in training even at very low bandwidths. We consider this to be a potential advantage in future real-world applications of RL where sample efficiency and rapid adaption to new environments are crucial.

It is noteworthy that deterministic (aka "greedy") experience selection seems to perform slightly better than stochastic experience selection, while in PER the opposite is generally the case [22]. We have two hypotheses for why this is the case. One, we note that in PER, the motivation for stochastic prioritization is to avoid low-error experiences never being sampled (nor re-prioritized) in many draws from the buffer. On the other hand, in suPER we only ever consider each experience once. Thus, if in stochastic experience selection a high-error experience through random chance is not shared on this one opportunity, it will never be seen by other agents. In a sense, we may prefer deterministic experience selection in suPER for the same reason we prefer stochastic selection in PER, which is to avoid missing out on potentially valuable experiences. Two, in all our current experiments we used (stochastic) PER when sampling training batches from the replay buffer of each agent. When using stochastic suPER, each experience therefore must pass through two sampling steps before being shared and learned on by another agent. It is possible that this dilutes the probability of a high-error experience being seen too much.

We also want to point out a slight subtlety in our theoretical motivation for suPER: We use the sending agent's td-error as a proxy for the usefulness of an experience for the receiving agent. We believe that this is justified in symmetric settings, and our experimental results support this. However, for future work, we would be interested to see if, first, using the td-error of each (potential) receiving agent could extend this approach to asymmetric settings, and if it could further improve performance even in symmetric settings. While this would effectively be a centralized-training approach, if it showed similar performance benefits as we have seen in symmetric settings for suPER, it could nevertheless be a promising avenue for further work. Secondly, we would be interested to see if the sending agent could learn to approximate each receiver's td-error locally, and thus retain the decentralized-with-communication training capability of our current approach. However, given that td-error is intrinsically linked to current policy and thus highly non-stationary, we expect there would be significant practical challenges to this approach.

Finally, we note that while we mainly focus on the DQN family of algorithms in this paper, suPER could be applied to other off-policy algorithm as well. If the improvements we see in this work hold for other algorithms and domains as well, this could improve multi-agent RL performance in many settings.

References

[1] Sanjeevan Ahilan and Peter Dayan. Correcting experience replay for multi-agent communication. *arXiv preprint arXiv:2010.01192*, 2020.

[2] Filippos Christianos, Lukas Schäfer, and Stefano Albrecht. Shared experience actor-critic for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 33:10707–10717, 2020.

[3] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems*, 29, 2016.

[4] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

[5] Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. Cooperative multi-agent control using deep reinforcement learning. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 66–83. Springer, 2017.

[6] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.

[7] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 33(6):750–797, 2019.

[8] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-second AAAI conference on artificial intelligence*, 2018.

[9] Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado Van Hasselt, and David Silver. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.

[10] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, DJ Strouse, Joel Z Leibo, and Nando De Freitas. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International conference on machine learning*, pages 3040–3049. PMLR, 2019.

[11] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. Multi-agent cooperation and the emergence of (natural) language. *arXiv preprint arXiv:1612.07182*, 2016.

[12] Joel Z. Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent reinforcement learning in sequential social dilemmas. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, AAMAS '17, page 464–473, Richland, SC, 2017. International Foundation for Autonomous Agents and Multiagent Systems.

[13] Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Ken Goldberg, Joseph E. Gonzalez, Michael I. Jordan, and Ion Stoica. RLlib: Abstractions for distributed reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2018.

[14] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

[15] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30, 2017.

[16] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. In *NIPS Deep Learning Workshop*. 2013.

[17] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[18] Peng Peng, Ying Wen, Yaodong Yang, Quan Yuan, Zhenkun Tang, Haitao Long, and Jun Wang. Multiagent bidirectionally-coordinated nets: Emergence of human-level coordination in learning to play starcraft combat games. *arXiv preprint arXiv:1703.10069*, 2017.

[19] Emanuele Pesce and Giovanni Montana. Improving coordination in multi-agent deep reinforcement learning through memory-driven communication. *CoRR*, abs/1901.03887, 2019.

[20] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International conference on machine learning*, pages 4295–4304. PMLR, 2018.

[21] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob N Foerster, and Shimon Whiteson. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research*, 21, 2020.

[22] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

[23] Jürgen Schmidhuber, Jieyu Zhao, and Nicol N Schraudolph. Reinforcement learning with self-modifying policies. In *Learning to learn*, pages 293–309. Springer, 1998.

[24] Sainbayar Sukhbaatar, Rob Fergus, et al. Learning multiagent communication with backpropagation. *Advances in neural information processing systems*, 29, 2016.

[25] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[26] Csaba Szepesvári. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.

[27] Ardi Tampuu, Tambet Matiisen, Dorian Kodelja, Ilya Kuzovkin, Kristjan Korjus, Juhan Aru, Jaan Aru, and Raul Vicente. Multiagent cooperation and competition with deep reinforcement learning. *PloS one*, 12(4):e0172395, 2017.

[28] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993.

[29] J. K Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sulivan, Luis Santos, Rodrigo Perez, Caroline Horsch, Clemens Dieffendahl, Niall L Williams, Yashas Lokesh, Ryan Sullivan, and Praveen Ravi. Pettingzoo: Gym for multi-agent reinforcement learning. *arXiv preprint arXiv:2009.14471*, 2020.

[30] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

[31] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pages 1995–2003. PMLR, 2016.

[32] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3):279–292, 1992.

[33] Annie Xie, Dylan P Losey, Ryan Tolsma, Chelsea Finn, and Dorsa Sadigh. Learning latent representations to influence multi-agent interaction. *arXiv preprint arXiv:2011.06619*, 2020.

[34] Lianmin Zheng, Jiacheng Yang, Han Cai, Weinan Zhang, Jun Wang, and Yong Yu. Magent: A many-agent reinforcement learning platform for artificial collective intelligence. *CoRR*, abs/1712.00600, 2017.

[35] Changxi Zhu, Mehdi Dastani, and Shihan Wang. A survey of multi-agent reinforcement learning with communication. *arXiv preprint arXiv:2203.08975*, 2022.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Additional Analysis of Bandwidth Sensitivity

We present here a more detailed analysis of bandwidth sensitivity of suPER-DQN in the three experience selection modes we discuss in the main text. Figure 4 shows the mean performance across five seeds for gaussian (left), quantile (middle) and stochastic (right) experience selection, at 1-2M timesteps (top) and at 250k timesteps (bottom). We can see that at 1-2M timesteps and a target bandwidth of 0.1, all three experience selection criteria perform similary. One thing that stands out is that stochastic selection has much lower performance at other target bandwidths, and also much less performance uplift compared to no-sharing DQN at 250k timesteps at any bandwidth. Gaussian experience selection appears to be less sensitive to target bandwidth, but upon closer analysis we found that it also was much less responsive in terms of how much actual bandwidth it used at different settings. Figure 5 (left) shows the actual bandwidth used by each selection criterion at different target bandwidths. We can see that quantile and stochastic experience hit their target bandwidth very well in

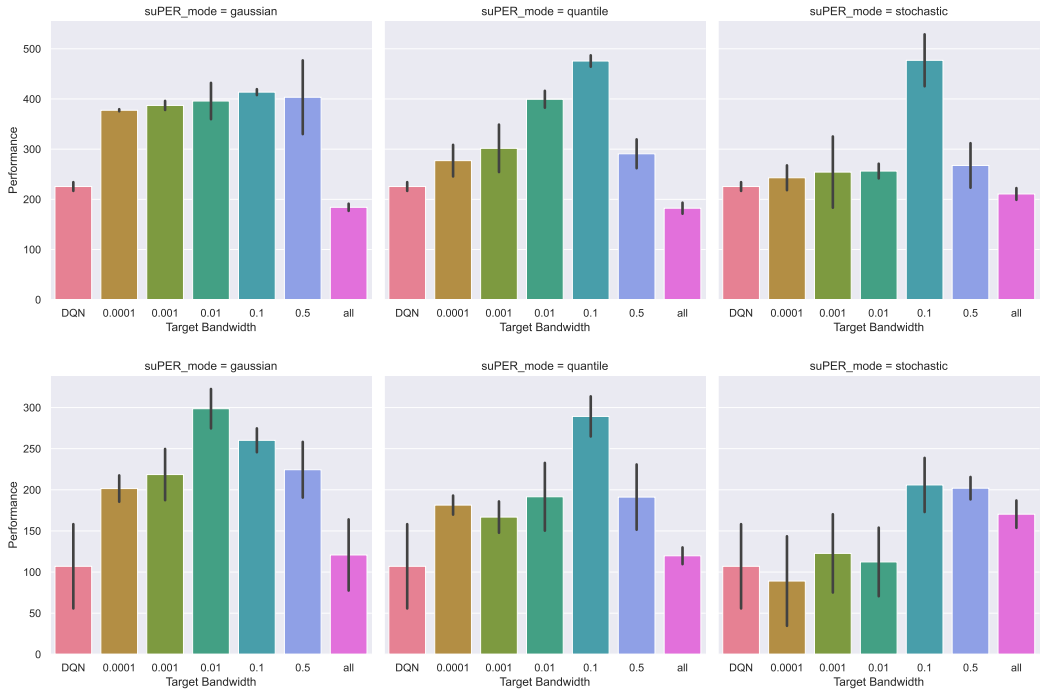

Figure 4: Performance of suPER with different experience selection and varying bandwidth in Pursuit at 1-2M timesteps (top) and at 250k timesteps (bottom).

general.[1] What stands out, however, is that gaussian selection vastly overshoots the target bandwidth at lower settings, never going significantly below 0.01 actual bandwidth.

What is a fairer comparison therefore is to look at performance versus actual bandwidth used for each of the approaches, which we do in Figure 5 (middle, at 1-2M timesteps, and right, at 250k timesteps). For these figures, we did the following: First, for each experience selection approach and target bandwidth, we computed the mean performance and mean actual bandwidth across the five seeds. Then, for each experience selection mode, we plotted these $(\text{meanactualbandwidth}, \text{meanperformance})$ (one for each target bandwidth) in a line plot.[2] The result gives us a rough estimate of how each approach's performance varies with actual bandwidth used. We see again that stochastic selection shows worse performance than quantile at low bandwidths, and early in training. We also see that gaussian selection very closely approximates quantile selection. Notice that gaussian selection never hits an exact actual bandwidth of 0.1, and so we cannot tell from these data if it would match quantile selection's performance at its peak. However, we can see that at the actual bandwidths that gaussian selection does hit, it shows very similar performance to quantile selection. As stated in the main text, our interpretation of this is that using mean and variance to approximate the exact distribution of absolute td-errors is a reasonable approximation, but that we might need to be more clever in selecting c in equation 3.

B Experiment Hyperparameters & Details

We performed all experiments using the open-source library **RLlib** [13], version 1.13.0, available under Apache license. Environments used are from PettingZoo [29], including SISL [5] and MAgent [34]. The suPER algorithm was implemented by modifying RLlib's standard DQN algorithm to

[1] Quantile selection overshoots at 1e-4 (0.0001) target bandwidth and is closer to 1e-3 actual bandwidth usage, which we attribute to a rounding error, as we ran these experiments with a window size of 1500 (1.5e+3), and a quantile of less than a single element is not well-defined.

[2] Because each data point now has variance in both x- and y-directions, it is not possible to draw error bars for these.

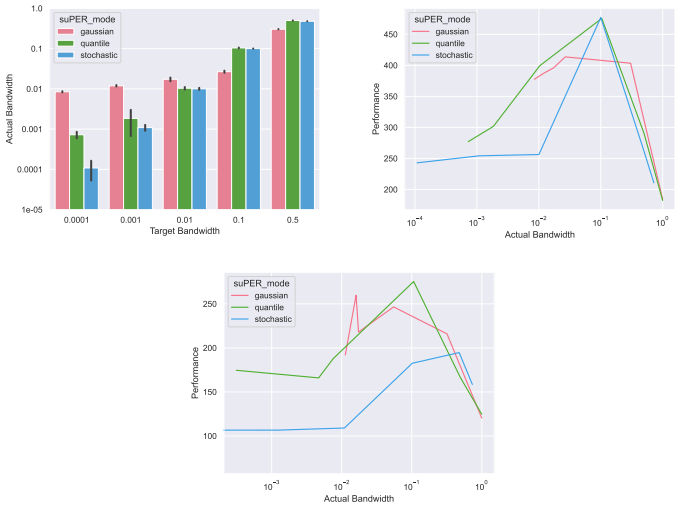

Figure 5: Left: Actual bandwidth used (fraction of experiences shared) at different target bandwidths. Middle, right: Performance compared to actual bandwidth used at 1-2M and 250k timesteps.

perform the suPER experience sharing between rollout and training. Table 1 lists all the DQN hyperparameters and environment settings we used for all the experiments. Experiments in the "Stability across hyperparameters" section had hyperparameters set to those listed in Table 1 except those specified in Figure 3. Any parameters not listed were left at their default values from RLlib 1.13.0, in particular we used dueling DDQN. Hyperparameters were tuned using a grid search; some of the combinations tested are also discussed in the "Stability across hyperparameters" section.

All experiments were repeated with five seeds, except exploration-configuration experiments in the "Stability across hyperparameters" section, which were repeated with three seeds. All plots show the mean and standard deviation of these seeds at each point in training. For technical reasons, individual experiment runs did not always report data at identical intervals. For instance, one run might report data when it had sampled 51000 environment timesteps, and another run might report at 53000 environment timesteps. In order to still be able to report a meaningful mean and standard deviation across repeated runs, we rounded down the timesteps reported to the next 10000 steps, i.e. we would take both the data above to represent each run's performance at 50000 steps. Where a run reported more than once in a 10000 step interval, we took the mean of its reports to represent that run's performance in the interval. Mean and standard deviation were calculated across this mean performance for each of the five seeds.

C Implementation & Reproducibility

All source code is included in the supplementary material and will be made available on publication under an open-source license. We refer the reader to the included README file, which contains instructions to recreate the experiments discussed in this paper.

Table 1: Hyperparameter Configuration Table

SISL: Pursuit			
Learning Parameters		Environment Parameters	
HyperParameters	**Value**	**HyperParameters**	**Value**
learning rate	0.00016	max cycles	500
batch size	32	x/y sizes	16/16
framework	torch	shared reward	False
prioritized replay_alpha	0.6	num evadors	30
prioritized replay eps	1e-06	obs range	7
horizon	500	n catch	2
dueling	True	freeze evaders	False
target network update_freq	1000	tag reward	0.01
buffer size	120000	constrained window	1.0
rollout fragment length	4	catch rewards	5
initial exploration epsilon	0.1	urgency reward	-0.1
final exploration epsilon	0	surrounded	True
		num agents(pursuers)	8
MAgent: Battle			
learning rate	1e-4	minimap mode	False
batch size	32	step reward	-0.005
framework	torch	dead penalty	-0.1
prioritized replay_alpha	0.6	attack penalty	-0.1
prioritized replay eps	1e-06	attack opponent reward	0.2
horizon	1000	max cycles	1000
dueling	True	extra features	False
target network update_freq	1200	map size	18
rollout fragment length	5		
buffer size	90000		
initial exploration epsilon	0.1		
final exploration epsilon	0		
MAgent: Adversarial Pursuit			
learning rate	1e-4	minimap mode	False
batch size	32	tag penalty	-0.2
framework	torch	max cycles	500
prioritized replay alpha	0.6	extra features	False
prioritized replay eps	1e-06	map size	18
horizon	500		
dueling	True		
target network update_freq	1200		
buffer size	90000		
rollout fragment length	5		
initial exploration epsilon	0.1		
final exploration epsilon	0		