

Samples Sharing and Prioritization in Multi-Agent Deep Reinforcement Learning

Thibault Lahire

Dassault Aviation

Saint-Cloud, France

thibault.lahire@dassault-aviation.com

ABSTRACT

By reusing experiences collected from past different policies, experience replay significantly improves the training efficiency of reinforcement learning algorithms. Rapid convergence occurs when learning is based on pertinent experiences that offer valuable information. Nonetheless, how to effectively combine experience replay with multi-agent reinforcement learning is still an open challenge. We study how sharing collected experiences helps the training process and show that sharing a small amount of selected experiences between agents improves the learning process compared to the baseline where each agent is independent. The shared experiences are selected by each agent on internal statistics, ensuring their meaningfulness. Our first results on the multi-agent Pursuit environment highlight an improvement by a large margin and need to be consolidated by complementary experiences.

KEYWORDS

Multi-Agent Deep Reinforcement Learning, Importance Sampling, Replay Buffers

ACM Reference Format:

Thibault Lahire. 2024. Samples Sharing and Prioritization in Multi-Agent Deep Reinforcement Learning. In *OptLearnMAS*, , 2024, IFAAMAS, 5 pages.

1 INTRODUCTION

Multi-Agent Systems [32] have benefited from Reinforcement Learning [29, RL] as it enabled to address many issues [36]. In particular, RL has made improvements in adaptive decision-making [3], handles partial observability [24], promotes emergent behavior and self-organization [16], provides decentralized control tools [25] and allows agents to generalize over different tasks thanks to transfer learning [8]. However, when each agent learns independently, Multi-Agent Reinforcement Learning (MARL) experiences difficulties in the learning process due to the non-stationarity of the environment. Even though the convergence guarantees of MARL are an active field of research [13], MARL still succeeds in learning in complex environments [30].

In deep RL, neural network policies and value functions can be learnt thanks to stochastic gradient descent algorithms [26, SGD] sampling an experience replay memory [18]. This replay memory, or replay buffer, stores the transitions encountered along the interaction with the environment. SGD-based algorithms exploit such buffers to learn relevant functions, such as the Q-function in the case of Deep Q-Networks [21], the return distribution for distributional

approaches [2], or an actor and a critic [11, 17] in the case of continuous state-action space problems. Most deep RL algorithms boil down to a sequence of SGD-based, supervised learning problems. In supervised learning, importance sampling can be used to speed up the convergence of SGD, by sampling non-uniformly the training set to reduce the variance of the stochastic gradient estimate. As there is a link between supervised learning and RL, accelerating the convergence thanks to non-uniform sampling has also been explored in the latter with Prioritized Experience Replay [27, PER], drawing inspiration from Prioritized Sweeping [22]. Extensions, modifications, and foundations of PER have been proposed, such as [34] and [15]. In this paper, we show that some of the techniques initially designed to accelerate convergence speed in the single-agent case can be used to help agents in their learning process.

In a partially observable environment which is "anonymous" (the environment behaves the same way for all agents) to homogeneous agents collecting at each time step their own reward, this work proposes a sharing-experience scheme among agents. Each agent is initialized with its own neural network and replay buffer. At each learning step, each agent updates its neural network parameters thanks to an SGD step with a mini-batch composed of experiences from its own replay buffer, as well as experiences collected by other agents. The relevance of a given experience to a given agent is checked before being used for the SGD step thanks to statistics specific to the updated agent. Our experience-sharing method compares to independent learning, where agents are initialized with their own neural network and replay buffer and do not share anything. We evaluate the benefit of sharing a small amount of experiences collected by other agents in a mixed collaborative-competitive environment.

This work is structured as follows. Section 2 covers related work in MARL and importance sampling for RL. Then Section 3 proposes an algorithm to improve learning in multi-agent settings by sampling efficiently useful experiences collected by the other agents. Section 4 empirically evaluates the proposed algorithm. We discuss each separate aspect of sampling, its explanations and perspectives. Section 5 discusses the limitations of this work and contrasts our findings with the existing literature. Section 6 summarizes and concludes.

2 BACKGROUND

We model our MARL problem as a Partially Observable Stochastic Game (POSG), a generalization of Stochastic (Markov) Games [19] to settings where agents are only able to observe parts of the state of the environment [12]. A POSG is a tuple $(M, \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{O}, \mathcal{T}_e)$, where M is the number of agents, \mathcal{S} is the set of all possible global states of the environment, $\mathcal{A} := A_1 \times A_2 \times \dots \times A_M$ represents the

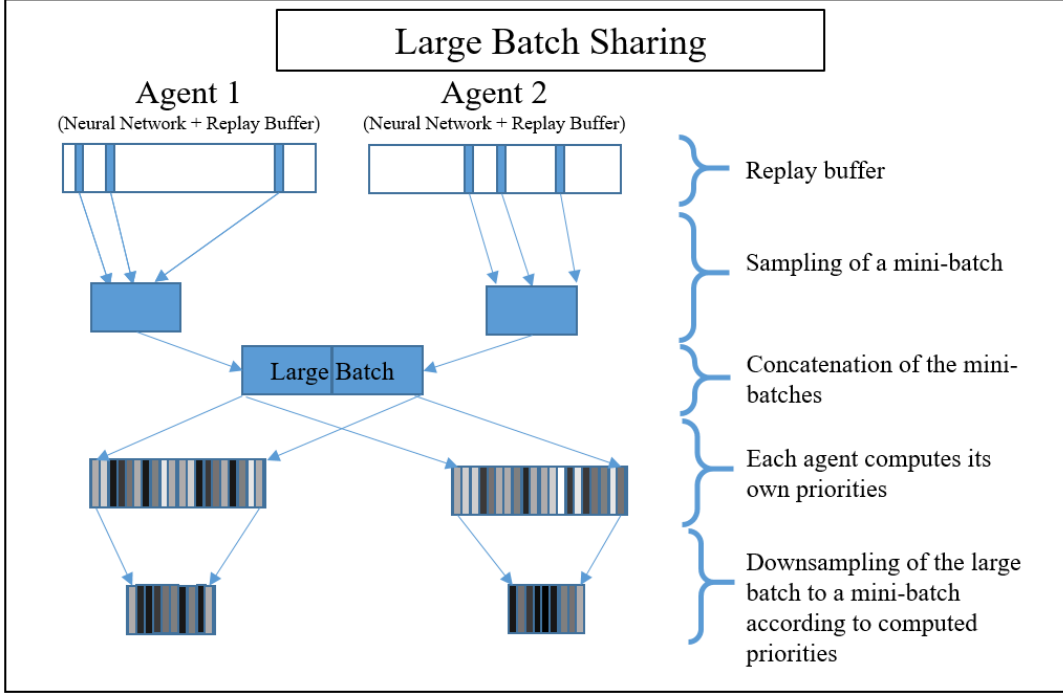


Figure 1: Illustration of the Large Batch Sharing algorithm in the case of two agents.

set of actions, where A_i is the action space of agent i , $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition probability function between global states, based on the joint action of the agents, $\mathcal{R} := R_1 \times R_2 \times \dots \times R_M$ is the reward function, where $R_i : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the individual reward function of agent i , $\mathcal{O} := O_1 \times O_2 \times \dots \times O_M$ is the set of observations with O_i representing the individual observation set for each agent i , and $\mathcal{T}_e : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{O})$ is the observation function.

The behaviour of an agent is defined by its policy $\pi_i : O_i \times A_i \rightarrow [0, 1]$. The performance of a policy can be assessed through its Q-function $Q^{\pi_i} = \mathbb{E}[\sum_t \gamma^t r_{i,t} | \pi]$, where $\pi = (\pi_1, \dots, \pi_M)$ is the joint policy of the agents acting in the environment, $\gamma \in [0, 1]$ is the discount factor, and $r_{i,t} = R_i(s_t, \mathbf{a}_t, s_{t+1})$ is the reward obtained by agent i at time step t for the joint action $\mathbf{a}_t \in \mathcal{A}$ at state $s_t \in \mathcal{S}$ and transitioning to the next state $s_{t+1} \in \mathcal{S}$. The goal of each agent is to find their policy π_i^* that has the largest possible Q-function.

As set previously, this work focuses on agents receiving their own reward in a partially observable environment that reacts to each agent the same way. If an action taken by an agent in a particular state yields reward, so does the environment for any other agent (the environment is said to be "anonymous"). We also make the assumption of homogeneous agents, sharing the same set of individual actions. All assumptions related to our multi-agent setting will be discussed in Section 5 along with extensions and limitations of our work.

In such a framework, two baseline algorithms can be used. First, one single agent, taking as input the observation and outputting the action for the agent receiving the observation can be trained thanks to the gathering of the agents' interactions with the environment. This method belongs to the Parameter Sharing (PS)

algorithms, where neural network(s) and replay buffer (if any) are shared among all agents. The second baseline we consider are independent learners, where each agent trains its own neural network(s) thanks to its own interactions with the environment. Neural network(s) parameters (and replay buffers) remain private.

Our contribution, consisting in designing an efficient sharing of experiences between agents, can be applied to any underlying model-free deep RL algorithm using a replay buffer. For pedagogical purposes, we focus on the Deep Q-Networks [21, DQN] algorithm, which is off-policy and designed for discrete actions, and we will explain in Section 5 how the proposed methods can be extended to other settings.

In single-agent RL with full observability, the optimal Q-function obeys equation $Q^*(s, a) = \mathbb{E}_{s', r} [r + \gamma \max_{a'} Q^*(s', a')]$, called the Bellman optimality equation. DQN is the approximate Value Iteration algorithm that uses a replay buffer of N samples (s, a, r, s') , a deep neural network Q_θ , and a few steps of gradient descent to minimize the Bellman optimality equation. Specifically, at each training step, DQN aims to take a gradient step on the empirical loss $\frac{1}{N} \sum_{i=1}^N \ell(Q_\theta(s_i, a_i), y_i)$, with $y_i = r_i + \gamma \max_{a'} Q_n(s'_i, a')$. Minimization of this empirical loss by SGD implies drawing at each step a mini-batch of B transitions from the replay buffer and taking a descent step $\theta_{t+1} = \theta_t - \eta d$ in the direction of the gradient estimate $d = \frac{1}{B} \sum_{i=1}^B \nabla_{\theta} \ell(Q_\theta(s_i, a_i), y_i)$, with learning rate η .

To the best of our knowledge, Prioritized Experience Replay [27, PER] is the first work introducing a non-uniform sampling of the replay buffer to accelerate the convergence of the DQN algorithm from an empirical perspective. At each iteration, PER samples a mini-batch according to the probability distribution induced by a

list of priorities, performs a gradient step and updates the priorities of the selected samples. PER drew inspiration from Prioritized Sweeping [22] and used TD errors as priorities without mathematical justification. Even though PER lacks theoretical foundations, a large number of publications experimentally demonstrate its benefits. Notably, the ablation study of [14] showed PER to be one of the most critical improvements over DQN. [15] cast prioritization of samples in RL as importance sampling and proposed non-uniform sampling schemes aiming at accelerating the convergence speed of SGD. It brought theoretical foundations to many prioritization algorithms, especially PER, and proposed more grounded sampling schemes. One of the proposed algorithm, LaBER (for Large Batch Experience Replay), consists in sampling uniformly a large batch from the replay buffer (hence the name), computing TD errors on this large batch and finally down-sampling the large batch to a mini-batch according to the distribution induced by the TD errors. Note that, contrarily to PER, the TD errors are up-to-date, and this difference allows a more accurate sampling scheme.

3 LARGE BATCH SHARING

In this work, we reuse the idea of a large batch, but we fill it with experiences collected by all agents which communicate. Then, we compute the TD errors on this large batch and we down-sample the large batch to a mini-batch thanks to the distribution induced by the TD errors. Let B be the mini-batch size, and N be the size of the replay buffers. Let K be the large batch size, so that $K = BM$, where M is the number of agents. Our algorithm, named Large Batch Sharing and illustrated in Figure 1, goes as follows.

For each agent $i \in [1; M]$, 1/ draw a mini-batch of size B from the replay buffer with distribution p_1 . Aggregate the $M - 1$ other mini-batches to obtain the large batch. 2/ On this large batch, compute probability distribution p_2 and down-sample the large batch to a mini-batch thanks to this distribution. 3/ Finally, perform the SGD step thanks to the samples in the mini-batch. We now detail the subtleties at each step previously described.

Step 1 of the algorithm could simply consist in drawing uniformly a mini-batch from the replay buffer (hence $p_1 = u$, where $u_i = 1/N, \forall i \in [1; N]$), but we will test different sampling, such as PER.

Step 2 computes the priority which determines which samples will yield the better SGD step thanks to distribution p_2 . Note that, if p_2 and p_1 are both uniform distributions, then our algorithm is strictly equivalent to a parameter sharing baseline where replay buffers are fully shared, but neural networks are not. We will, of course, test this baseline but our aim is to explore different sampling schemes. In particular, as proposed in LaBER, we take p_2 as the distribution induced by the TD-errors. By doing this, we ensure that the samples coming from other agents really help in taking better SGD steps.

Step 3 must be carefully implemented. As explained previously, the SGD step in its standard form for the DQN algorithm, sampling uniformly a replay buffer with a mini-batch of size B is

$$\theta_{t+1} = \theta_t - \eta \frac{1}{B} \sum_{i=1}^B \nabla_{\theta} l_i(\theta) \quad (1)$$

with $l_i(\theta) = (r_i + \gamma \max_{a'_i} Q_{\hat{\theta}}(s'_i, a'_i) - Q_{\theta}(s_i, a_i))^2$. When importance sampling is used to obtain the mini-batch, the update equation




Figure 2: Performance of the four agents in Pursuit at 800k timesteps.

depends on the probability distribution q (over the N items of the replay buffer) used for sampling (see [15] for more details):

$$\theta_{t+1} = \theta_t - \eta \frac{1}{B} \sum_{i=1}^B \frac{1}{Nq_i} \nabla_{\theta} l_i(\theta). \quad (2)$$

From a theoretical point of view, we cannot discriminate between the two. Indeed, in our multi-agent framework, the data distribution on which the learning is rooted is unknown, contrarily to supervised learning. The distribution of collected experiences from which agents can learn optimally their tasks remains unknown. Hence we cannot assess the quality of an update equation in our multi-agent RL case using importance sampling. We will test different update equations.

4 EXPERIMENTS

Even though many experiences remain to be run, the first results obtain on the SISL (Stanford Intelligent Systems Laboratory) environment named Pursuit appear promising [10]. In the Pursuit scenario, a mixed collaborative-competitive environment is presented, involving a team of pursuers aiming to capture a group of evaders within a grid-world containing obstacles. The evaders, depicted in blue, move randomly, while the pursuers, represented in red, are under the control of RL agents. When a group of two or more agents successfully surrounds an evader, each agent receives a reward, and the evader is eliminated from the environment. The episode concludes either when all evaders are captured or after 500 steps, whichever comes first. Pursuers earn a small reward for being adjacent to an evader (even without complete surrounding) and incur a slight negative reward per timestep, encouraging them to complete episodes promptly. The setup involves 8 pursuers and 30 evaders. We perform the training for 800k timesteps and average our results over 10 seeds.

As explained earlier, the two baselines are 1/ the independent DQN agents (IDQN), and 2/ the parameter sharing agent (PS) owning one neural network and replay buffer collecting all experiments. Our agent uses p_1 as the uniform distribution, and p_2 as the distribution induced by the up-to-date TD errors. We use Eq. 1 for the

Table 1: Hyperparameters

<i>Environment hyperparameters</i>	
max cycles	500
shared reward	False
horizon	500
surrounded	True
tag reward	0.01
constrained window	1.0
obs range	7
x/y sizes	16/16
num evaders	30
n catch	2
n agents (pursuers)	8
urgency reward	-0.1
catch rewards	5
<i>CNN network hyperparameters</i>	
CNN layers	[32, 64, 64]
Stride	1
Kernel size	[2, 2]
<i>DQN hyperparameters</i>	
learning rate	0.00016
mini-batch size	32
buffer size	120000
initial exploration epsilon	0.1
final exploration epsilon	0.001
target network update freq	1000
factor K	8

SGD update. The fourth algorithm tested (IDQN same RB) consists in taking p_1 and p_2 as the uniform distribution. In this case, each agent has its own neural network but the replay buffer is shared by all agents. All hyperparameters for the DQN base agent are the same for all the tested configurations to ensure a fair comparison and are reported in Table 1.

Figure 1 shows the average sum of rewards per episode at the end of training for each tested algorithm. This result advocates for a high benefit of our sampling but remains to be consolidated by other experiments. In particular, a sensitivity analysis with respect to hyper-parameters specific to our method have to be tested. We plan to analyze the impact of: 1/ the amount of experiences to share between agents, 2/ the sampling distribution used for p_1 and p_2 , and 3/ the update equation for SGD (either Eq. 1 or Eq. 2).

5 DISCUSSIONS, LIMITATIONS AND RELATED WORK

Our work restricts to homogeneous agents. An experience can be useful to another agent only if this agent can also take the action of the experience. Moreover, in the case of heterogeneous agents, where agents do not share the same action space, the baseline algorithms are different. In particular, the parameter sharing baseline is not applicable. Our work also restricts to "anonymous" environments and could not be extended to a different setting, where a

certain action combined with a certain observation does not yield the same reward to all agents. Note that this kind of environment is rare in practice, and when encountered, it is often in the case of heterogeneous agents, as in [4].

Our contribution is illustrated on collaborative systems, and it also applies to mixed and competitive settings, as long as agents remain homogeneous and the environment anonymous. We found more intuitive to highlight the benefits of sharing experiences on collaborative settings, as parameter sharing is rarely seen in practice on competitive environments. Indeed, in such settings, agents often explore the environment for their own profit, avoiding sharing information to maintain an advantage in the environment knowledge compared to the competitors.

In cooperative multi-agent systems, where all agents share the same reward at each step, our contribution can be used as well. However, note that cooperative tasks are harder to solve compared to collaborative ones, where each agent earns its own reward, due to the credit assignment problem. Even though it remains to be experimentally verified, we suspect our method to be less beneficial in this setting. Sharing experiences between agents is likely to make the credit assignment even more harder to solve, since knowing property of one experiment might help the credit assignment.

This work focuses on the DQN algorithm, as it is a common practice in the single-agent literature to study importance sampling and prioritization on this base agent. Our contribution extends directly to all off-policy algorithms using a replay buffer to learn a value function. This encompasses algorithms designed for continuous actions, such as DDPG [17], TD3 [9], SAC [11], and also distributional RL agents, such as C51 [2], QR-DQN [6], and IQN [5]. Sharing experiences can also be performed with on-policy algorithms such as PPO [28]. However, as prioritizing samples with policy-based algorithms remains a difficult problem, even in the single-agent setting, we keep this extension of our algorithm for future works.

This enumeration of limitations is necessary to understand the scope of our study and to place it in the literature. Perhaps the most similar work to ours is [35]. The authors use PER to select the shared experiences, but do not check the relevance of a particular sample for the agent which receives it, which is what we propose in our work. [7] couple multi-task RL algorithms with a task-sampling policy based on the intrinsic motivation paradigm. [23] propose an exploration strategy which enables to fill the replay buffer with experiences collected from a distribution supposed to be beneficial for the learning process.

Prioritized sampling has also been studied in multi-agent RL in settings different to ours. In cooperative tasks, an extended literature has been produced in recent years studying the interplay between non-stationarity, credit assignment, and importance sampling. For instance, [20] derives an importance sampling scheme allowing to correct the distribution of the replay buffers, hence enabling a better convergence. [1] compare to [20] but frame their study in model-based MARL. [31] study how PER can improve off-line MARL thanks to a trajectory selection using Graph Attention Networks [33].

6 CONCLUSION

This work focuses on an efficient experience sharing scheme between homogeneous agents learning in an anonymous environment. This scheme draws inspiration from the optimization literature, and our work remains an empirical study as no convergence or improvement guarantees can be derived. This article remains a work in progress, and the first experimental results highlight the benefit of sharing a small amount of experiences between agents.

In the near future, complementary experiments have to be run. Different environments have to be tested, as well as different numbers of agents. In particular, we would like to confirm / infirm the following intuition: Our method brings a significant improvement on tasks where the independent learners baseline performs well, in environments necessitating specialization of agents. Complementary experiments will also study from an empirical point of view the amount of experiments to be shared between agents yielding the best learning improvement.

7 ACKNOWLEDGMENTS

This work highlights some mathematical issues at stake in multi-agent deep reinforcement learning and is to be pursued in the AI illustrations carried out in the EICACS (European Initiative Collaborative Air Combat Standardisation) project.

This publication was co-funded by the European Union. Its contents are the sole responsibility of the author and do not necessarily reflect the views of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.



Co-funded by
the European Union

REFERENCES

- [1] Eugenio Bargiacchi, Timothy Verstraeten, Diederik M Roijers, and Ann Nowé. 2020. Model-based multi-agent reinforcement learning with cooperative prioritized sweeping. *arXiv preprint arXiv:2001.07527* (2020).
- [2] Marc G Bellemare, Will Dabney, and Rémi Munos. 2017. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning*. PMLR, 449–458.
- [3] Lucian Busoniu, Robert Babuska, and Bart De Schutter. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38, 2 (2008), 156–172.
- [4] Jeancarlo Arguello Calvo and Ivana Dusparic. 2018. Heterogeneous Multi-Agent Deep Reinforcement Learning for Traffic Lights Control. *AICS*, 2–13.
- [5] Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. 2018. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning*. PMLR, 1096–1105.
- [6] Will Dabney, Mark Rowland, Marc Bellemare, and Rémi Munos. 2018. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [7] Carlo D’Eramo and Georgia Chalvatzaki. 2022. Prioritized Sampling with Intrinsic Motivation in Multi-Task Reinforcement Learning. In *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [8] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems* 29 (2016).
- [9] Scott Fujimoto, Herke Hoof, and David Meger. 2018. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*. 1587–1596.
- [10] Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. 2017. Cooperative multi-agent control using deep reinforcement learning. In *Autonomous Agents and Multiagent Systems: AAMAS 2017 Workshops, Best Papers, São Paulo, Brazil, May 8–12, 2017, Revised Selected Papers* 16. Springer, 66–83.
- [11] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *International Conference on Machine Learning*. 1861–1870.
- [12] Eric A Hansen, Daniel S Bernstein, and Shlomo Zilberstein. 2004. Dynamic programming for partially observable stochastic games. In *Proceedings of the 19th national conference on Artificial intelligence*. 709–715.
- [13] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor. 2019. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 33, 6 (2019), 750–797.
- [14] Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. 2018. Rainbow: Combining Improvements in Deep Reinforcement Learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [15] Thibault Lahire, Matthieu Geist, and Emmanuel Rachelson. 2022. Large Batch Experience Replay. In *International Conference on Machine Learning*. PMLR, 11790–11813.
- [16] Zhengping Li, Cheng Hwee Sim, and Malcolm Yoke Hean Low. 2006. A survey of emergent behavior and its impacts in agent-based systems. In *2006 4th IEEE international conference on industrial informatics*. IEEE, 1295–1300.
- [17] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. In *ICLR (Poster)*.
- [18] Long-Ji Lin. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning* 8, 3-4 (1992), 293–321.
- [19] Michael L Littman. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*. Elsevier, 157–163.
- [20] Yongsheng Mei, Hanhan Zhou, Tian Lan, Guru Venkataramani, and Peng Wei. 2023. MAC-PO: Multi-agent experience replay via collective priority optimization. *arXiv preprint arXiv:2302.10418* (2023).
- [21] 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. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [22] Andrew W Moore and Christopher G Atkeson. 1993. Prioritized sweeping: Reinforcement learning with less data and less time. *Machine learning* 13, 1 (1993), 103–130.
- [23] Isack Thomas Nicholas and Dae-Ki Kang. 2022. Robust experience replay sampling for multi-agent reinforcement learning. *Pattern Recognition Letters* 155 (2022), 135–142.
- [24] Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P How, and John Vian. 2017. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *International Conference on Machine Learning*. PMLR, 2681–2690.
- [25] Liviu Panait and Sean Luke. 2005. Cooperative multi-agent learning: The state of the art. *Autonomous agents and multi-agent systems* 11 (2005), 387–434.
- [26] Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics* (1951), 400–407.
- [27] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2016. Prioritized Experience Replay. In *ICLR (Poster)*.
- [28] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [29] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [30] Ardi Tampuu, Tanel Matiisen, Dorian Kodelja, Ilya Kuzovkin, Kristjan Korjus, Juhan Aru, Jaan Aru, and Raul Vicente. 2017. Multiagent cooperation and competition with deep reinforcement learning. *PLoS one* 12, 4 (2017), e0172395.
- [31] Qi Tian, Kun Kuang, Furui Liu, and Baoxiang Wang. 2023. Learning from good trajectories in offline multi-agent reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 11672–11680.
- [32] Wiebe Van der Hoek and Michael Wooldridge. 2008. Multi-agent systems. *Foundations of Artificial Intelligence* 3 (2008), 887–928.
- [33] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- [34] Che Wang and Keith Ross. 2019. Boosting soft actor-critic: Emphasizing recent experience without forgetting the past. *arXiv preprint arXiv:1906.04009* (2019).
- [35] Yishen Wang and Zongzhang Zhang. 2019. Experience selection in multi-agent deep reinforcement learning. In *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 864–870.
- [36] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. 2021. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of reinforcement learning and control* (2021), 321–384.