

Offline-to-Online Multi-Agent Reinforcement Learning with Offline Value Function Memory and Sequential Exploration

Hai Zhong
IIIS, Tsinghua University
Beijing, China

zhongh22@mails.tsinghua.edu.cn

Zhuoran Li
IIIS, Tsinghua University
Beijing, China
lizr20@mails.tsinghua.edu.cn

Xun Wang
IIIS, Tsinghua University
Beijing, China

wang-x24@mails.tsinghua.edu.cn

Longbo Huang*
IIIS, Tsinghua University
Beijing, China
longbohuang@tsinghua.edu.cn

ABSTRACT

Offline-to-Online Reinforcement Learning has emerged as a powerful paradigm, leveraging offline data for initialization and online fine-tuning to enhance both sample efficiency and performance. However, most existing research has focused on single-agent settings, with limited exploration of the multi-agent extension, i.e., Offline-to-Online Multi-Agent Reinforcement Learning (O2O MARL). In O2O MARL, two critical challenges become more prominent as the number of agents increases: (i) the risk of unlearning pre-trained Q-values due to distributional shifts during the transition from offline-to-online phases, and (ii) the difficulty of efficient exploration in the large joint state-action space. To tackle these challenges, we propose a novel O2O MARL framework called **Offline Value Function Memory with Sequential Exploration (OVMSE)**. First, we introduce the **Offline Value Function Memory (OVM)** mechanism to compute target Q-values, preserving knowledge gained during offline training, ensuring smoother transitions, and enabling efficient fine-tuning. Second, we propose a decentralized **Sequential Exploration (SE)** strategy tailored for O2O MARL, which effectively utilizes the pre-trained offline policy for exploration, thereby significantly reducing the joint state-action space to be explored. Extensive experiments on the StarCraft Multi-Agent Challenge (SMAC) demonstrate that OVMSE significantly outperforms existing baselines, achieving superior sample efficiency and overall performance.

KEYWORDS

Multi-Agent Reinforcement Learning; Offline-to-Online Reinforcement Learning

ACM Reference Format:

Hai Zhong, Xun Wang, Zhuoran Li, and Longbo Huang*. 2025. Offline-to-Online Multi-Agent Reinforcement Learning with Offline Value Function

Memory and Sequential Exploration. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

1 INTRODUCTION

Multi-Agent Reinforcement Learning (MARL) has achieved remarkable success across various domains, including mastering complex video games [1, 3, 28], optimizing warehouse logistics [25], enabling robotic soccer [17], and performing bi-manual dexterous manipulation [6]. However, these successes often come at the cost of low sample efficiency and high computational overhead, as MARL algorithms must explore a joint state-action space that grows exponentially with the number of agents. A promising approach to alleviate this computational burden is Offline-to-Online (O2O) Reinforcement Learning (RL). In recent years, O2O RL has achieved significant progress [10, 14, 15, 19, 29, 32–34], leveraging advances in Offline RL by utilizing offline datasets [7, 8] to provide strong initial policies and pre-trained value functions. By learning a high-quality pre-trained policy from offline data, single-agent O2O RL can significantly reduce the need for extensive exploration and further enhance performance through interaction during the online phase. However, existing results primarily focus on single-agent scenarios, and the important O2O MARL setting has received only very limited attention, e.g., [12].

O2O MARL faces two major challenges. First, exploration during the transition from the offline-to-online phase introduces a distributional shift, that can result in significant *unlearning* of the pre-trained Q-values in the early stages of online learning. To demonstrate this phenomenon, we conducted an experiment analyzing how pre-trained Q-values evolve during online fine-tuning. We collected trajectories for the offline pre-trained agent interacting with the environment and stored them into a replay buffer. During the online phase, we evaluated how online value estimations for these samples evolved. As shown in Figure 1, Multi-agent Conservative Q-Learning (MACQL) [13], Multi-agent Cal-QL (MACal-QL) [19], and the action proposal method from [9] exhibit clear signs of unlearning, as the fine-tuned Q-values exhibit a rapid drop during the initial stage of online learning. This suggests that the algorithms have forgotten the previously learned optimal actions. This unlearning behavior hinders the efficiency of online fine-tuning, as the policy

*Corresponding to Longbo Huang (longbohuang@tsinghua.edu.cn).



This work is licensed under a Creative Commons Attribution International 4.0 License.

has to relearn knowledge that was already acquired during the offline phase. In contrast, our proposed algorithm OVMSE preserves the offline knowledge and achieves fast online fine-tuning.

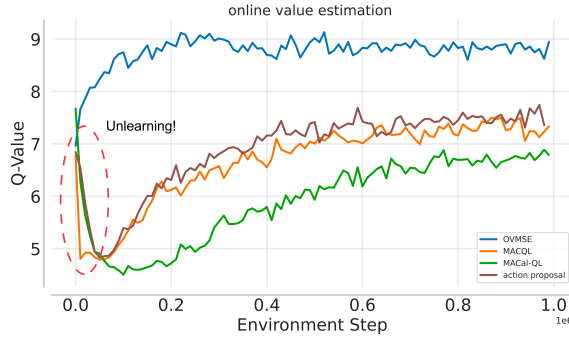


Figure 1: Q-value evolution for different algorithms during the online phase, starting from pre-trained Q-value from offline learning. OVMSE clearly preserves the offline Q-value, while other methods experience rapid unlearning (i.e., online value estimations for these samples decrease rapidly).

Second, efficient exploration is critical for O2O MARL. Exploration at the initial stage of online learning can lead to overestimation of Q-values for unseen state-action pairs. While this overestimation is necessary for agents to experiment with new actions and improve upon the offline pre-trained policy, inefficient exploration may result in assigning high Q-values to sub-optimal actions, requiring significant trial and error to correct. This issue is particularly pronounced in multi-agent systems, where the joint state-action space grows exponentially with the number of agents. However, unlike training from scratch, O2O MARL benefits from a pre-trained offline policy, which provides a much stronger starting point than a randomly initialized policy. Therefore, an effective exploration strategy for O2O MARL needs to focus on exploring more efficiently within a reduced joint state-action space, rather than exhaustively searching through the entire space.

To address these two challenges, we propose a novel O2O MARL algorithm named **Offline Value Function Memory and Sequential Exploration (OVMSE)**. Specifically, OVMSE consists of two key components. The first component, **OVM**, introduces a novel way of computing target Q-values and solves the unlearning problem. With OVM, even when the offline pre-trained Q-values degrade due to distributional shifts during the early stages of online learning, the algorithm allows for the quick recovery of the offline pre-trained values. The second key component, **Sequential Exploration**, is designed to enable efficient exploration after transitioning to the online phase. Inspired by the sequential update mechanism [4, 30], SE restricts exploration to one agent at a time by allowing a single agent to select a random action while the others follow their respective policies. SE reduces the complexity of the joint state-action space and improves the quality of exploration, enabling more efficient fine-tuning and better performance. We further develop

decentralized SE for decentralized execution. We empirically validate OVMSE on a collection of easy, hard, and super-hard tasks in SMAC. Our results demonstrate superior performance compared to baseline methods.

The paper is organized as follows. Section 2 discusses related literature, Section 3 introduces the background knowledge for MARL and QMIX [23], Section 4 describes the proposed OVMSE algorithm, and Section 5 presents the experimental results and discussions. The main contributions of our paper are summarized as follows:

- (i) **Key Challenges in O2O MARL:** We identify and analyze two key challenges in O2O MARL: (1) Offline-learned Q-values can be unlearned during the initial stage of online learning, leading to the degradation of offline pre-trained policies; (2) Inefficient exploration in the exponentially large joint state-action space can result in slow online fine-tuning.
- (ii) **OVMSE Algorithm:** We propose OVMSE, a novel framework that combines Offline Value Function Memory (OVM) and Sequential Exploration (SE) to address both the unlearning and exploration issues. OVMSE is designed to facilitate robust and efficient online fine-tuning, reducing the sample complexity while achieving superior performance.
- (iii) **Extensive Empirical Evaluation on SMAC:** We conduct extensive empirical evaluations of OVMSE on a range of tasks—spanning easy, hard, and super-hard difficulties—in the StarCraft Multi-Agent Challenge (SMAC). Our results demonstrate that OVMSE significantly outperforms baseline methods. OVMSE shows significantly improved online performance, higher sample efficiency, and faster fine-tuning, underscoring its practical applicability to complex multi-agent environments.

2 RELATED WORK

O2O RL and O2O MARL. Many previous works on O2O RL focus on reusing offline data during online training [2, 10, 14]. [32] proposes retaining a copy of the offline policy and using a hybrid of the offline and online fine-tuned policies during online training, thereby preserving the performance achieved during the offline phase. A related approach to our work is presented in [9] which utilizes the offline policy to suggest actions for online target value estimation. In the context of O2O MARL, relatively little research attention has been given. [12] investigates O2O MARL by using offline data during the online phase and proposes metrics for sampling from offline data. In addition to our proposed value function memory, we employ a multi-agent sequential exploration strategy inspired by multi-agent sequential rollout [4, 30] to reduce the exploration space in cooperative multi-agent systems. This improves online learning efficiency, as the offline pre-trained policy can focus on targeted exploration rather than an exhaustive random search of the action space, which is typically required when training from scratch.

Offline MARL. The principles used to address accumulated extrapolation error in offline single-agent RL also apply to offline MARL. A common strategy is to incorporate pessimism towards out-of-distribution (OOD) states and actions, which has proven effective in the offline MARL setting as well. Both [21] and [24] extend

the Conservative Q-Learning (CQL) framework [13] to the multi-agent domain. [21] identifies that value functions in MARL are more prone to local optima as the number of agents increases and proposes using zeroth-order optimization to avoid being trapped in these local optima. [24] demonstrates that independently penalizing each agent's OOD actions can reduce the overall conservatism in value estimation, leading to improved empirical performance. [31] takes a different approach by ensuring that only actions present in the dataset are used for target Q-value estimation, thereby promoting conservatism for OOD states and actions. Another promising direction is to treat offline MARL problems as supervised learning tasks. For example, [27] and [18] frame offline MARL as sequence modeling problems, leveraging the power of decision transformers [5] to achieve strong empirical performance. Furthermore, [16] and [35] employ diffusion models as policy backbones, demonstrating good performance and the ability to learn diverse strategies.

3 PRELIMINARIES

3.1 Multi-Agent Reinforcement Learning Formulation

A cooperative MARL problem can be described as a decentralized partially observable Markov decision process (Dec-POMDP) [20], which can be represented as a tuple $\langle \mathcal{N}, \mathcal{S}, \mathcal{O}, \mathcal{A}, P, r, \gamma, T \rangle$.

- $\mathcal{N} = \{1, 2, \dots, N\}$ is the set of agents.
- \mathcal{S} is the set of states of the environment.
- $\mathcal{O} = \mathcal{O}_1 \times \mathcal{O}_2 \times \dots \times \mathcal{O}_N$ is the joint observation space, where \mathcal{O}_i is the set of observations available to agent i .
- $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_N$ is the joint action space, where \mathcal{A}_i is the set of actions available to agent i .
- Given the current state $s \in \mathcal{S}$ and the joint action $\mathbf{a} \in \mathcal{A}$, $P(s' | s, \mathbf{a})$ and $r(s, \mathbf{a})$ describe the probability of transitioning from s to the next state s' , and the immediate reward shared by all agents, respectively.
- $\gamma \in [0, 1]$ is the discount factor, and T is the time horizon of the process.

At time step t , each agent maintains its observation-action history $\tau_i \in \mathcal{T}_i = \{(o_{i,0}, a_{i,0}, \dots, o_{i,t}, a_{i,t})\}$, and the joint observation-action history is $\tau = (\tau_1, \tau_2, \dots, \tau_n)$. The objective is to find a policy for each agent $\pi_i : \mathcal{T}_i \rightarrow \mathcal{A}_i$ that maximizes the expected discounted return

$$\mathbb{E} \left[\sum_{t=0}^T \gamma^t \cdot r(s_t, \mathbf{a}_t) \right],$$

where s_t and \mathbf{a}_t are the state and the joint action at time step t , respectively.

3.2 QMIX

In this section, we present QMIX [22, 23], which is a popular MARL algorithm following the Centralized Training with Decentralized Execution (CTDE) paradigm. In our work, we utilize QMIX as the backbone for both the offline and online phases.

Under the CTDE framework, each agent has an individual action-value function $Q_i(\tau_i, a_i)$, based on its observation-action history. Additionally, QMIX introduces a joint action-value function $Q_{\text{tot}}(\tau, \mathbf{a})$, constructed using a mixing network f_s , which combines individual

agent action-values according to:

$$Q_{\text{tot}}(\tau, \mathbf{a}) = f_s(Q_1(\tau_1, a_1), \dots, Q_n(\tau_n, a_n)). \quad (1)$$

The mixing network satisfies the monotonicity constraint,

$$\frac{\partial Q_{\text{tot}}}{\partial Q_i} \geq 0 \quad \forall i \in \mathcal{N}, \quad (2)$$

thereby maintaining that the joint action-value function Q_{tot} is a non-decreasing function with respect to each individual action-value function Q_i . Hence, QMIX satisfies the Individual-Global-Maximum (IGM) property [26], which guarantees that the joint optimal action can be constructed by combining each agent's individual optimal actions based on their local observations:

$$\arg \max_{\mathbf{a}} Q_{\text{tot}}(\tau, \mathbf{a}) = \left(\arg \max_{a_1} Q_1(\tau_1, a_1), \dots, \arg \max_{a_n} Q_n(\tau_n, a_n) \right). \quad (3)$$

4 OFFLINE VALUE FUNCTION MEMORY AND SEQUENTIAL EXPLORATION (OVMSE)

We now introduce our Offline Value Function Memory and Sequential Exploration (OVMSE) algorithm for O2O MARL. First, we present our online training mechanism, OVM. Next, we describe the designed decentralized sequential exploration mechanism, SE. Finally, we outline our offline training procedure.

4.1 Offline Value Function Memory

In the online phase, we introduce our Offline Value Function Memory (OVM) target. We describe the functionality of our proposed OVM training objective for online training. After offline training, we retain a copy of the pre-trained target value function, denoted as $\bar{Q}_{\text{tot-offline}}$. We introduce the OVM target, defined as:

$$\bar{Q}_{\text{OVM}} = \max \left(\bar{Q}_{\text{tot-offline}}(\tau, \mathbf{a}), r + \gamma \max_{\mathbf{a}'} \bar{Q}_{\text{tot}}(\tau', \mathbf{a}') \right). \quad (4)$$

where \bar{Q}_{tot} is the online target Q-function, τ' and \mathbf{a}' are the joint observation-action history and joint action at the next time step, $r + \gamma \max_{\mathbf{a}'} \bar{Q}_{\text{tot}}(\tau', \mathbf{a}')$ is the online temporal difference target, and $\bar{Q}_{\text{tot-offline}}(\tau, \mathbf{a})$ represents the offline memory of the agents with joint history τ and joint action \mathbf{a} . The OVM target selects the maximum value between the offline value function memory and the online temporal difference target.

During the online phase, the value function is trained to minimize the mean squared error (MSE) between its estimation and the OVM target, as well as the MSE between its estimation and the online temporal difference target, i.e.,

$$\begin{aligned} \mathcal{L}_{\text{OVM}} = & (1 - \lambda_{\text{memory}}) \left(Q_{\text{tot}}(\tau, \mathbf{a}) - \left(r + \gamma \max_{\mathbf{a}'} \bar{Q}_{\text{tot}}(\tau', \mathbf{a}') \right) \right)^2 \\ & + \lambda_{\text{memory}} \left(Q_{\text{tot}}(\tau, \mathbf{a}) - \bar{Q}_{\text{OVM}} \right)^2, \end{aligned} \quad (5)$$

which introduces a trade-off between the OVM target and the online temporal difference target. Here, λ_{memory} is a memory coefficient that controls the balance between these two objectives.

We further propose an annealing schedule for λ_{memory} , which gradually adjusts the influence of the offline value memory over

time. The annealing schedule is defined as follows:

$$\lambda_{\text{memory}} = \min \left(\lambda_{\text{memory_end}}, \left(1 - \left(\frac{1 - \lambda_{\text{memory_end}}}{T} \right) \cdot t \right) \right) \quad (6)$$

where $\lambda_{\text{memory_end}}$ specifies the final value for λ_{memory} , T is the annealing duration, and t represents the current environment step. The motivation for gradually decreasing λ_{memory} is to maintain the memory of the offline value memory during the transition from offline-to-online learning. As the agent interacts with the environment and stores experiences in the replay buffer, the offline values are progressively adjusted towards the online temporal difference target values, facilitating further improvement.

This training objective ensures that the online value function preserves the pre-trained offline values while allowing the agent to follow the online temporal difference target when it is larger. As a result, the agent retains knowledge from the offline phase while effectively exploring new strategies and improving its performance during the online phase. The intuition behind OVM is as follows: the Q-values evolve based on the pre-trained offline Q-values. For unseen actions that were assigned low Q-values during offline training, their values can increase as they are explored online, allowing the agent to discover and exploit better strategies while also retaining offline knowledge.

4.2 Multi-agent Sequential Exploration

ϵ -greedy exploration is one of the most widely used strategies for exploration in multi-agent reinforcement learning (MARL) algorithms [11, 23]. In this strategy, each agent independently selects a random action with probability ϵ_t , which gradually decreases over environment steps:

$$\epsilon_t = \min \left(\epsilon_{\text{end}}, \left(1 - \left(\frac{\epsilon_{\text{start}} - \epsilon_{\text{end}}}{T} \right) \cdot t \right) \right), \quad (7)$$

where ϵ_{start} and ϵ_{end} specify the start and terminal values for ϵ_t , T is the annealing duration, and t is the current environment step. During the initial online phase, when ϵ is high, multiple agents are likely to explore randomly and simultaneously, resulting in exploration that resembles a random search across the exponentially large joint state-action space. However, this approach is inefficient for O2O MARL, where a pre-trained offline policy already provides a strong starting point. Instead of exhaustively searching the entire joint state-action space, exploration should focus on a smaller, more targeted subset to refine the policy rather than starting from scratch.

Inspired by the multi-agent sequential update scheme [4, 30], which sequentially updates different agents' policies, we introduce the concept of sequential exploration for O2O MARL. Specifically, we propose a sequential ϵ -greedy exploration strategy. At each time step, agents collectively decide whether to explore based on the probability ϵ_t , which decays with the number of environment steps. If exploration is chosen, only one randomly selected agent performs a random action, while all other agents act greedily, following the current policy. This coordinated approach allows the system to efficiently leverage the pre-trained policy while facilitating targeted exploration to further refine and improve the policy during the online phase.

Algorithm 1 Offline Value Function Memory with Sequential Exploration (OVMSE)

```

1: Offline pretrain via a CQL-based algorithm and retain  $\bar{Q}_{\text{OVM}}$ .
2: Initialize:  $\lambda_{\text{memory}} \leftarrow 1.0$ ,  $\epsilon \leftarrow \epsilon_{\text{start}}$ 
3: for each environment step  $t$  do
4:   Explore the environment via SE (Algorithm 2).
5:   Draw samples from the offline dataset and online buffer.
6:   Calculate  $\mathcal{L}_{\text{OVM}}$  according to (5).
7:   Update agents via QMIX.
8:   Update  $\lambda_{\text{memory}}$  and  $\epsilon_t$  according to (6) and (7).
9: end for

```

Algorithm 2 Decentralized Multi-Agent Sequential ϵ -Greedy Exploration

```

1: Input: exploration rate  $\epsilon_t$  at environment step  $t$ 
2: Set decentralized exploration probability:  $\epsilon_{\text{dec}_t} = \epsilon_t / N$ 
3: for each environment step  $t$  do
4:   for each agent  $i \in \{1, 2, \dots, N\}$  do
5:     (Agents operate in parallel)
6:     Sample  $u \sim \text{Uniform}(0, 1)$ 
7:     if  $u < \epsilon_{\text{dec}_t}$  then
8:       Agent  $i$  selects a random action
9:     else
10:      Agent  $i$  selects a greedy action
11:   end if
12: end for
13: end for

```

Decentralized Sequential Exploration. However, this coordinated exploration requires communication between agents during execution, which may not always be feasible in decentralized execution settings. To address this limitation, we introduce a decentralized version of sequential exploration with a simple modification. We define a decentralized exploration probability, $\epsilon_{\text{dec}_t} = \epsilon_t / N$, where N is the number of agents. Each agent independently decides to explore with probability ϵ_{dec_t} , ensuring that, on average, the number of agents exploring at any given time matches that of the centralized version. This decentralized approach maintains the benefits of sequential exploration while ensuring compatibility with decentralized execution.

4.3 Offline Training for OVMSE

In this section, we introduce the offline training procedure for OVMSE. During offline training, we utilize QMIX and Conservative Q-Learning (CQL) as our backbone algorithms. First, we minimize the squared temporal difference error as part of our offline training objective:

$$\mathcal{L}_{\text{QMIX}} = \left[Q_{\text{tot}}(\tau, \mathbf{a}) - \left(r + \gamma \max_{\mathbf{a}'} \bar{Q}_{\text{tot}}(\tau', \mathbf{a}') \right) \right]^2, \quad (8)$$

where τ' and \mathbf{a}' denote the next observation-action history and joint action, respectively, and \bar{Q}_{tot} represents the target joint action-value function.

In addition to minimizing the temporal difference error, we incorporate the CQL training objective into the offline phase. CQL-based

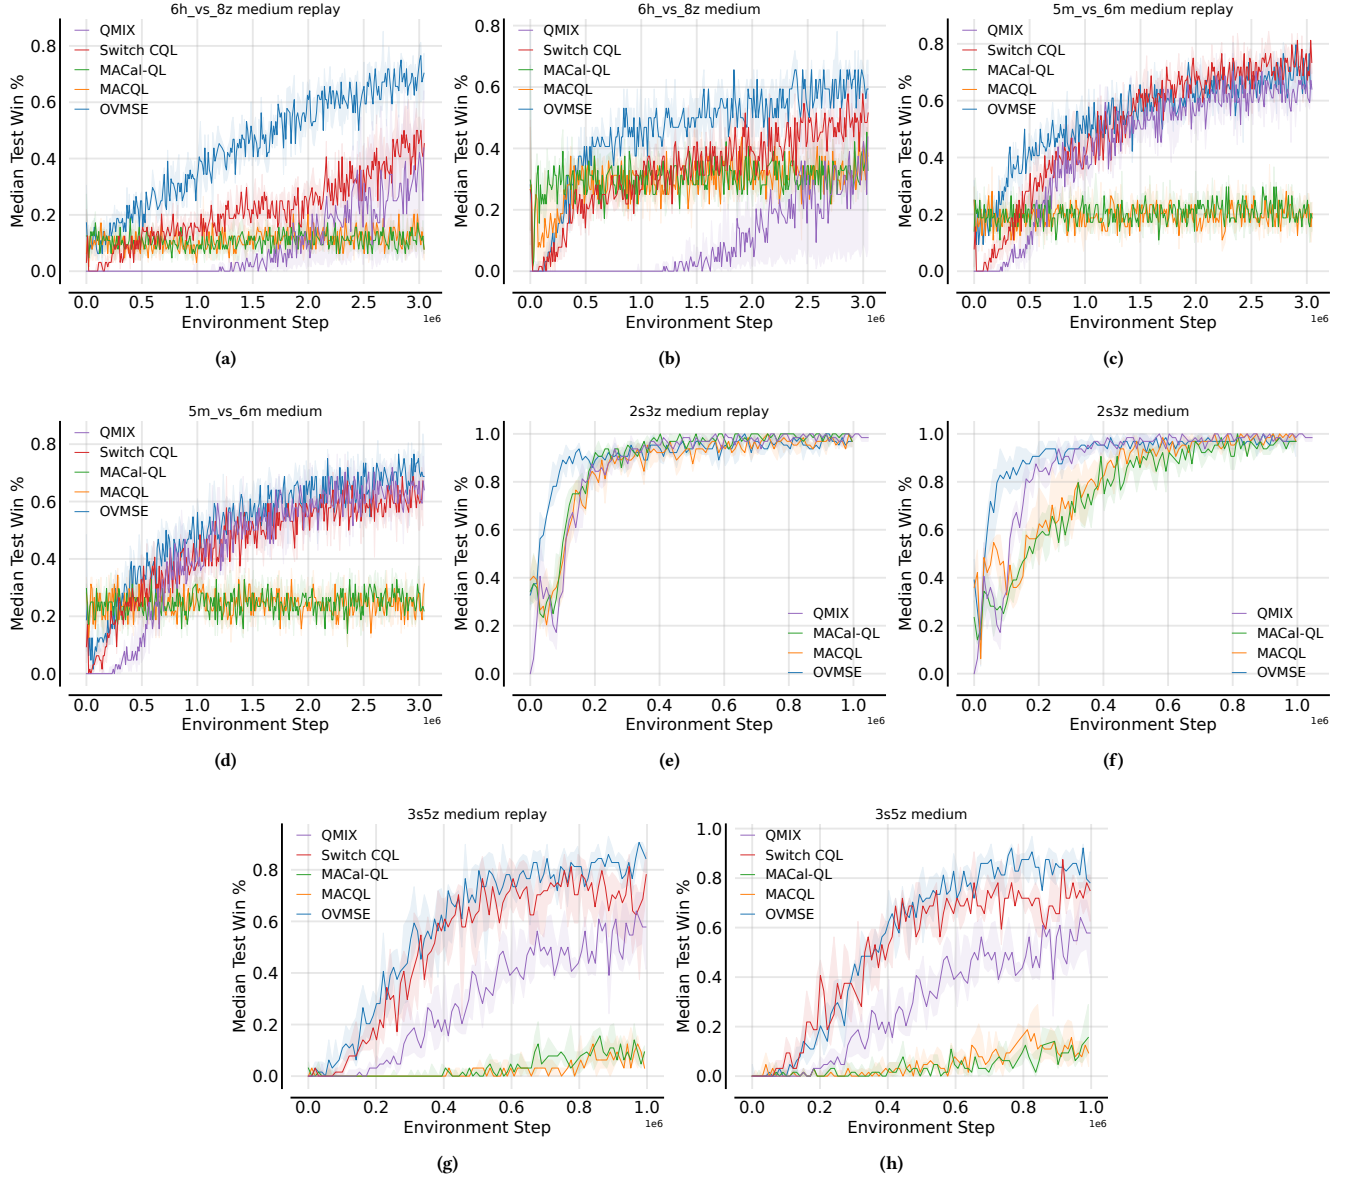


Figure 2: Median test win rates with respect to environment steps for OVMSE and baselines across different tasks, pre-trained with medium or medium replay datasets. Our results demonstrate that OVMSE preserves offline knowledge, enables efficient exploration, and leads to faster fine-tuning and superior overall performance.

methods introduce a regularization term to the Q-value function’s training objective, defined as:

$$\mathcal{L}_{CQL} = \mathbb{E}_{\tau \sim D, a \sim \mu} [Q(\tau, a)] - \mathbb{E}_{\tau \sim D, a \sim D} [Q(\tau, a)], \quad (9)$$

where μ denotes the sampling distribution (e.g., sampling from the current policy or a uniform distribution). This regularization term intuitively penalizes the Q-values of actions not present in the dataset while promoting the Q-values of actions within the dataset.

The overall offline training objective combines the QMIX temporal difference error and the CQL regularization:

$$\mathcal{L}_{\text{offline}} = \mathcal{L}_{\text{QMIX}} + \alpha \mathcal{L}_{CQL}, \quad (10)$$

where α is a hyperparameter that controls the strength of the CQL regularization.

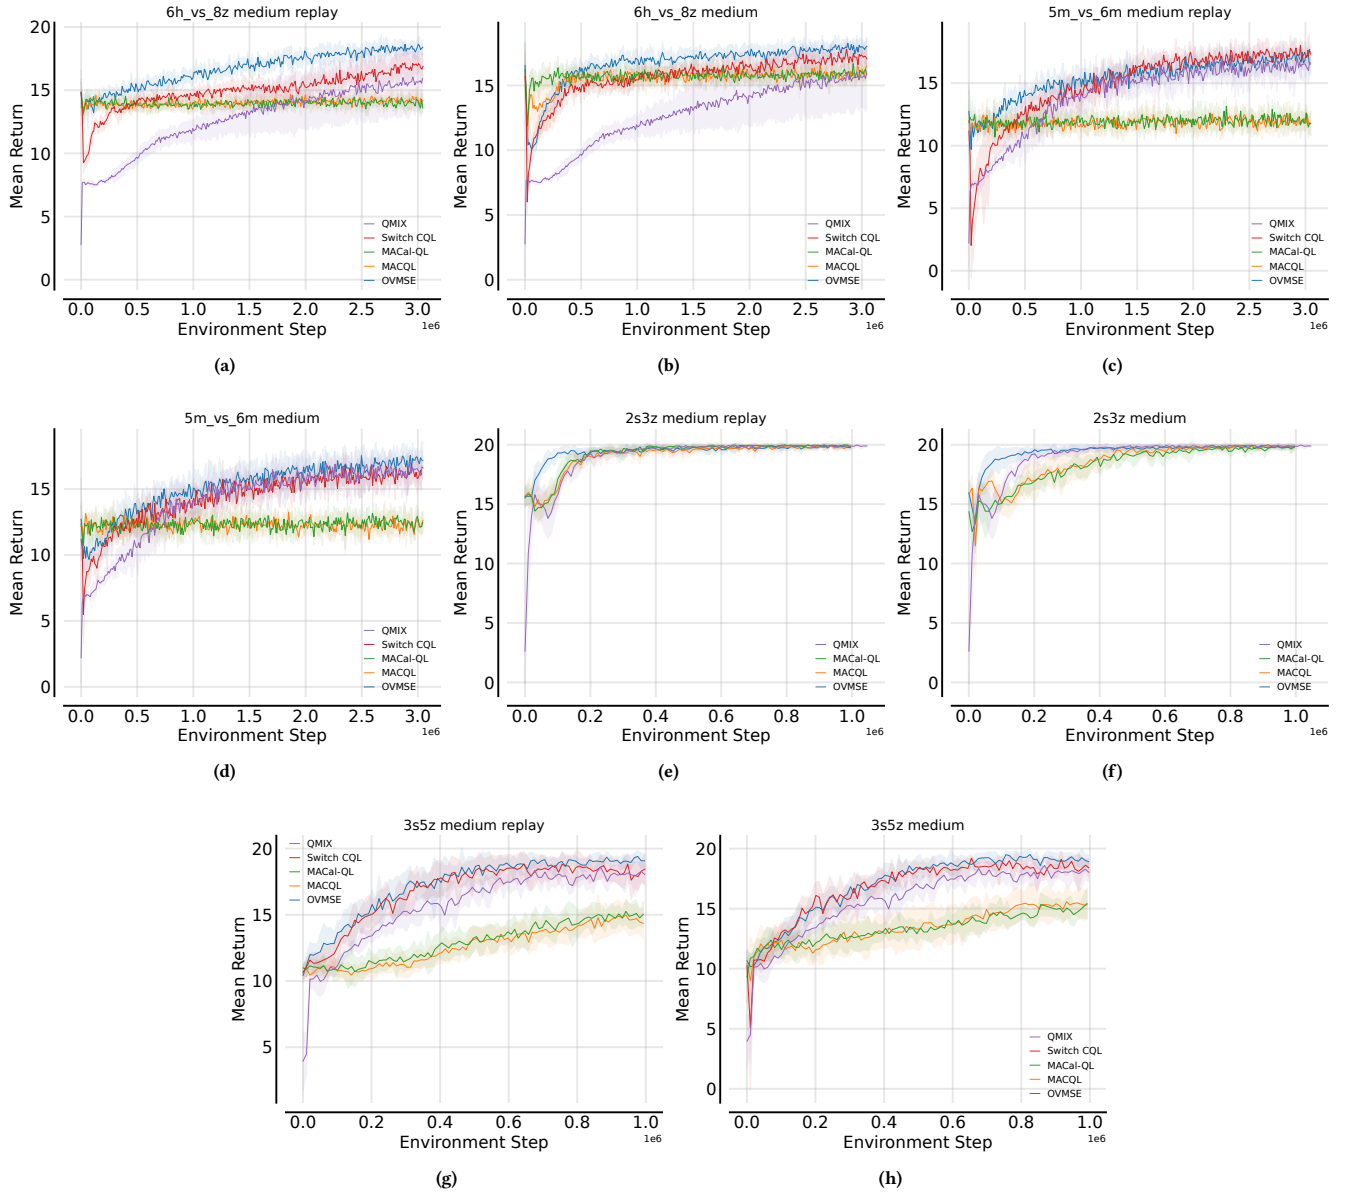


Figure 3: Mean test returns for all tasks under different algorithms.

5 EXPERIMENT

5.1 Setup

Codebase, Offline Datasets, and Experiment Setup. We evaluate our proposed OVMSE algorithm using the StarCraft Multi-Agent Challenge (SMAC), a key benchmark for MARL algorithms, across four tasks: *2s3z* (easy), *3s5z* (easy), *5m_vs_6m* (hard), and *6h_vs_8z* (super hard). For the tasks *2s3z*, *5m_vs_6m*, and *6h_vs_8z*, we use offline datasets provided by CFCQL [24]. We collected our own dataset for *3s5z*, as it is not publicly available. Offline training leverages medium and medium-replay datasets [8]. During the online phase, we apply the CQL loss for OVMSE in *2s3z*, as it stabilizes

online training. However, for *3s5z*, *5m_vs_6m*, and *6h_vs_8z*, we exclude the CQL loss because it hinders online learning by preventing the discovery of new strategies. Our implementation, including code and hyperparameters, is based on PYMARL2 [11] and CFCQL. Fine-tuning is conducted for 1 million environment steps in *2s3z* and *3s5z*, and for 3 million steps in *5m_vs_6m* and *6h_vs_8z*.

Baselines. We compare OVMSE against four baselines: (i) MACQL [13], which uses CQL loss throughout both the offline and online phases; (ii) MACal-QL, which applies Cal-QL (an algorithm designed for O2O single-agent RL) [19] to MARL; (iii) QMIX, trained from scratch [23]; and (iv) Switch CQL in *3s5z*, *5m_vs_6m*, and

6h_vs_8z, which switches off the CQL loss during the online phase. All baselines use QMIX as the backbone.

Mixing Ratio and Hyperparameters. Following [19], we introduce the concept of mixing ratio, a value ranging from 0 to 1, representing the percentage of offline data used during online training. For instance, a mixing ratio of 0 means no offline data is used during online training, while a ratio of 0.5 means half the data comes from offline and half from online interactions. In this work, we select the optimal mixing ratio hyperparameters for all methods from the set $\{0.0, 0.1, 0.3, 0.5\}$. Our experiments show that OVMSE performs best with a mixing ratio of 0.0 or 0.1. The detailed hyperparameter settings can be found in Appendix.

Results. We report the median test win rates over environment steps in Figure 2 and the mean test return in Figure 3. All results are averaged across 6 random seeds. Additionally, we conduct 640 tests for each seed after online fine-tuning, with the final median test win rates presented in Table 1.

5.2 Results

Figure 2 and 3 present the training curves across four different tasks, pre-trained using both the medium and medium replay datasets. Our results show that OVMSE enables fast online fine-tuning and achieves superior overall performance. We compare OVMSE against baseline methods across three key aspects: (i) online performance, (ii) sample efficiency, and (iii) performance drop during the offline-to-online transition.

Table 1 shows the final fine-tuned performance for all tasks (except for *2s3z*, where all algorithms achieve near 100% win rates). In the *6h_vs_8z* medium replay task, OVMSE outperforms QMIX and Switch-CQL by more than 20% in win rate. Similarly, in the *3s5z* medium tasks, OVMSE leads other baselines by roughly 10% in win rate, and in the *3s5z* medium replay, *6h_vs_8z* medium, and *5m_vs_6m* medium tasks, OVMSE demonstrates a lead of more than 5% in test win rate.

In terms of sample efficiency, OVMSE performs significantly better than the baselines. For instance, in the *6h_vs_8z* medium and medium replay tasks, OVMSE achieves a 40% win rate approximately 1.5 million environment steps ahead of other baselines. In the *5m_vs_6m* medium task, OVMSE achieves a 40% win rate about 0.5 million environment steps earlier than QMIX and Switch-CQL, with these baselines needing a total of 1.5 million additional steps to match OVMSE’s performance. In the *2s3z* medium and medium replay tasks, OVMSE achieves an 80% win rate with a lead of more than 100k environment steps over MACQL, MACal-QL, and QMIX.

Furthermore, OVMSE experiences significantly less performance drop compared to all baselines during the offline-to-online transition. Notably, although other baselines typically use a higher mixing ratio than OVMSE, OVMSE still exhibits significantly less performance drop in the *6h_vs_8z* medium, medium replay, *2s3z* medium and medium replay, and *5m_vs_6m* medium and medium replay tasks. This difference in performance drop is more evident in terms of mean test return, which is shown in Figure 3. These results show that our OVMSE algorithm efficiently utilizes offline data and avoids unlearning.

5.3 Ablation Study

In this section, we conduct additional experiments to demonstrate the effectiveness of both OVM and Sequential Exploration. Through ablation studies, we aim to answer two key questions:

- (i) Does OVM preserve offline knowledge and lead to better performance?
- (ii) Does SE enable efficient exploration and lead to better performance?
- (iii) Does OVMSE rely on reusing offline data during online phase?

We perform these ablation studies on the *2s3z* medium task, *6h_vs_8z* medium and medium replay tasks. As shown in Figure 4(a), (b), and (c), we set the mixing ratio to 0.0 for all tasks except for the *6h_vs_8z* medium task, where a mixing ratio of 0.1 is applied. Additionally, we conduct a specific ablation study on the *2s3z* medium task to examine the effect of reusing offline data during the online phase. As illustrated in Figure 4, we report the mean test returns as a function of environment steps, as this metric more effectively captures performance drops during the initial stages of online learning.

OVM is effective. Our results demonstrate that OVM plays a crucial role in O2O MARL. In the *2s3z* medium task, using SE alone without OVM results in a more significant performance drop during the initial stage of the online phase. Additionally, using OVM alone results in a much smaller performance drop compared to the action proposal approach [9]. In the *6h_vs_8z* medium replay task, OVMSE significantly outperforms SE alone, due to the preservation of offline knowledge by OVM. Similarly, in the *6h_vs_8z* medium task, OVMSE outperforms SE, with SE requiring 2 million steps to catch up.

SE is effective. Our results also demonstrate that SE is highly effective. In the *2s3z* medium task, OVMSE outperforms OVM alone by a large margin in the first 200k steps, due to SE’s more efficient exploration. In the *6h_vs_8z* medium task, using SE alone outperforms Switch CQL, where the only difference between the two algorithms is the use of ϵ -greedy exploration versus decentralized SE. This clearly demonstrates the effectiveness of SE.

OVMSE has minimal reliance on offline data during online training. We conduct an ablation study on *2s3z* medium to examine the effect of reusing offline data during online training for OVMSE. We report the test returns at 200k environment steps (instead of the final fine-tuned test returns, as different mixing ratio values all achieve an optimal final test return). As shown in Table 2, OVMSE with a mixing ratio of 0.0 and 0.1 achieves fast and stable online fine-tuning during the online phase, demonstrating that OVMSE is not sensitive to data distribution shift and effectively preserves offline learned Q-values. This highlights a clear advantage of OVMSE, as over-reliance on offline data can lead to slow online fine-tuning, especially when the offline data is of poor quality.

5.4 Discussions

In this section, we discuss the role of offline pre-trained knowledge in O2O MARL. Offline knowledge is beneficial in several ways. First, it helps guide online exploration much more efficiently than training from scratch. It also provides a strong initial policy. Furthermore, the offline pre-trained knowledge can facilitate faster

Table 1: The final median win rates for different tasks across algorithms, except for 2s3z in which all algorithms achieve near 100% win rate. The results are obtained based on 640 tests for each seed after the completion of online fine-tuning.

Task	OV MSE	Switch CQL	MACal-QL	MACQL	QMIX
6h_vs_8z medium replay	0.71250	0.50859	0.11641	0.12969	0.32031
6h_vs_8z medium	0.60078	0.53359	0.37187	0.34141	0.32031
5m_vs_6m medium	0.69531	0.63203	0.24688	0.24453	0.64063
5m_vs_6m medium replay	0.69297	0.71484	0.21641	0.20234	0.64063
3s5z medium	0.86719	0.74766	0.12969	0.17188	0.53359
3s5z medium replay	0.80625	0.74687	0.11797	0.04100	0.53359




Figure 4: We present the results of the ablation study and conclude that both OVM and SE are useful and effective.

Task	Mixing Ratio 0.0	Mixing Ratio 0.1	Mixing Ratio 0.3
2s3z medium	19.58	19.64	18.41

Table 2: We conduct an ablation study for OVMSE with different mixing ratios on 2s3z medium and provide the mean test return at 200k environment steps. The results show that OVMSE has a minimal reliance on offline data (performing best with a mixing ratio of 0.0 and 0.1), as it preserves the offline learned Q-values.

online learning and lead to better performance, as clearly demonstrated in the *6h_vs_8z* medium replay task shown in Figure 4 (b), where OVMSE outperforms SE by a large margin, thanks to its offline knowledge. However, we also observe that offline knowledge can hinder online learning in some cases. This issue likely arises because current offline training algorithms are not designed for the online phase. A good offline policy requires some degree of over-fitting to the offline datasets. However, for O2O MARL, this over-fitting can cause online learning to require many samples to correct or even result in the agent sticking to the offline policy, thereby preventing improvement. This observation motivated the introduction of the memory coefficient λ_{memory} and its corresponding annealing schedule, which allow OVMSE to gracefully adjust offline knowledge to the level of online returns.

6 CONCLUSIONS

In this work, we investigate the O2O MARL. We identify two major challenges for O2O MARL: (i) During the initial stage of online learning, distributional shift can degrade the offline pre-trained Q-values, leading to unlearning of previously acquired knowledge; (ii) Inefficient exploration in the exponentially large joint state-action space

can slow down online fine-tuning. To address these challenges, we introduce Offline Value Function Memory (OVM) to preserve offline knowledge and decentralized sequential exploration strategy for efficient online exploration. Our approach, OVMSE, enables stable and fast online fine-tuning and leads to improved fine-tuned results. Experiments across various tasks in SMAC demonstrate the superior performance of OVMSE.

ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China Grants 52450016 and 52494974.

REFERENCES

- [1] Bowen Baker, Ingmar Kanitscheider, Todor Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. 2020. Emergent Tool Use From Multi-Agent Autocurricula. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SkxpxJBKwS>
- [2] Philip J. Ball, Laura Smith, Ilya Kostrikov, and Sergey Levine. 2023. Efficient Online Reinforcement Learning with Offline Data. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 1577–1594. <https://proceedings.mlr.press/v202/ball23a.html>

- [3] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębniak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680* (2019).
- [4] Dimitri Bertsekas. 2021. Multiagent Reinforcement Learning: Rollout and Policy Iteration. *IEEE/CAA Journal of Automatica Sinica* 8, 2 (2021), 249–272. <https://doi.org/10.1109/JAS.2021.1003814>
- [5] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. 2021. Decision Transformer: Reinforcement Learning via Sequence Modeling. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 15084–15097. https://proceedings.neurips.cc/paper_files/paper/2021/file/7f489f642a0ddb10272b5c31057f0663-Paper.pdf
- [6] Yuanpei Chen, Yiran Geng, Fangwei Zhong, Jiaming Ji, Jiechuang Jiang, Zongqing Lu, Hao Dong, and Yaodong Yang. 2024. Bi-DexHands: Towards Human-Level Bimanual Dexterous Manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 5 (2024), 2804–2818. <https://doi.org/10.1109/TPAMI.2023.3339515>
- [7] Claude Formanek, Asad Jeewa, Jonathan Shock, and Arnun Pretorius. 2023. Off-the-Grid MARL: Datasets and Baselines for Offline Multi-Agent Reinforcement Learning. In *Extended Abstract at the 2023 International Conference on Autonomous Agents and Multiagent Systems*. AAMAS.
- [8] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4RL: Datasets for Deep Data-Driven Reinforcement Learning. *arXiv:2004.07219* [cs.LG]
- [9] Hengyuan Hu, Suvir Mirchandani, and Dorsa Sadigh. 2024. Imitation Bootstrapped Reinforcement Learning. In *Proceedings of Robotics: Science and Systems*. Delft, Netherlands. <https://doi.org/10.15607/RSS.2024.XX.056>
- [10] Hao Hu, Yiqin Yang, Jianing Ye, Chengjie Wu, Ziqing Mai, Yujing Hu, Tangjie Lv, Changjie Fan, Qianchuan Zhao, and Chongjie Zhang. 2024. Bayesian Design Principles for Offline-to-Online Reinforcement Learning. In *Forty-first International Conference on Machine Learning*. <https://openreview.net/forum?id=HLHQxMydFk>
- [11] Jian Hu, Siying Wang, Siyang Jiang, and Weixun Wang. 2023. Rethinking the Implementation Tricks and Monotonicity Constraint in Cooperative Multi-agent Reinforcement Learning. In *ICLR Blogposts 2023* (May 1, 2023). <https://iclr-blogposts.github.io/2023/blog/2023/riit/> <https://iclr-blogposts.github.io/2023/blog/2023/riit/>
- [12] Jiechuan Jiang and Zongqing Lu. 2023. Online Tuning for Offline Decentralized Multi-Agent Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 7 (Jun. 2023), 8050–8059. <https://doi.org/10.1609/aaai.v37i7.25973>
- [13] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative Q-Learning for Offline Reinforcement Learning. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1179–1191. https://proceedings.neurips.cc/paper_files/paper/2020/file/0d2b2061826a5df322116a5085a6052-Paper.pdf
- [14] Seunghyun Lee, Younggyo Seo, Kimin Lee, Pieter Abbeel, and Jinwoo Shin. 2021. Offline-to-Online Reinforcement Learning via Balanced Replay and Pessimistic Q-Ensemble. In *5th Annual Conference on Robot Learning*. <https://openreview.net/forum?id=AljXhEl6j5W>
- [15] Kun LEI, Zhengmao He, Chenhao Lu, Kaizhe Hu, Yang Gao, and Huazhe Xu. 2024. Uni-O4: Unifying Online and Offline Deep Reinforcement Learning with Multi-Step On-Policy Optimization. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=tbFBh3LMKi>
- [16] Zhuoran Li, Ling Pan, and Longbo Huang. 2023. Beyond Conservatism: Diffusion Policies in Offline Multi-agent Reinforcement Learning. *arXiv:2307.01472* [cs.AI] <https://arxiv.org/abs/2307.01472>
- [17] Siqi Liu, Guy Lever, Zhe Wang, Josh Merel, S. M. Ali Eslami, Daniel Hennes, Wojciech M. Czarnecki, Yuval Tassa, Shayegan Omidshafiei, Abbas Abdolmaleki, Noah Y. Siegel, Leonard Hasenclever, Luke Marris, Saran Tunyasuvunakool, H. Francis Song, Markus Wulfmeier, Paul Muller, Tuomas Haarnoja, Brendan Tracey, Karl Tuyls, Thore Graepel, and Nicolas Heess. 2022. From motor control to team play in simulated humanoid football. *Science Robotics* 7, 69 (2022), eabo0235. <https://doi.org/10.1126/scirobotics.abo0235> <https://www.science.org/doi/pdf/10.1126/scirobotics.abo0235>
- [18] Linghui Meng, Muning Wen, Chenyang Le, Xiyun Li, Dengpeng Xing, Weinan Zhang, Ying Wen, Haifeng Zhang, Jun Wang, Yaodong Yang, et al. 2023. Offline pre-trained multi-agent decision transformer. *Machine Intelligence Research* 20, 2 (2023), 233–248.
- [19] Mitsuhiro Nakamoto, Simon Zhai, Anikait Singh, Max Sobol Mark, Yi Ma, Chelsea Finn, Aviral Kumar, and Sergey Levine. 2023. Cal-QL: Calibrated Offline RL Pre-Training for Efficient Online Fine-Tuning. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 62244–62269. https://proceedings.neurips.cc/paper_files/paper/2023/file/c44a04289beaf0a7d968a94066a1d696-Paper-Conference.pdf
- [20] Frans A. Oliehoek and Christopher Amato. 2016. *A Concise Introduction to Decentralized POMDPs* (1st ed.). Springer Publishing Company, Incorporated.
- [21] Ling Pan, Longbo Huang, Tengyu Ma, and Huazhe Xu. 2022. Plan Better Amid Conservatism: Offline Multi-Agent Reinforcement Learning with Actor Rectification. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 17221–17237. <https://proceedings.mlr.press/v162/pan22a.html>
- [22] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *Journal of Machine Learning Research* 21, 178 (2020), 1–51. <http://jmlr.org/papers/v21/20-081.html>
- [23] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 4295–4304. <https://proceedings.mlr.press/v80/rashid18a.html>
- [24] Jianzhun Shao, Yun Qu, Chen Chen, Hongchang Zhang, and Xiangyang Ji. 2023. Counterfactual Conservative Q Learning for Offline Multi-agent Reinforcement Learning. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 77290–77312. https://proceedings.neurips.cc/paper_files/paper/2023/file/f3f2ff9579ba6deeb89ca2fe1f0b99c-Paper-Conference.pdf
- [25] Yi Shen, Benjamin McClosky, and Michael Zavlanos. 2023. Multi-agent reinforcement learning for resource allocation in large-scale robotic warehouse sortation centers. In *IEEE CDC 2023*. <https://www.amazon.science/publications/multi-agent-reinforcement-learning-for-resource-allocation-in-large-scale-robotic-warehouse-sortation-centers>
- [26] Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International conference on machine learning*. PMLR, 5887–5896.
- [27] Wei-Cheng Tseng, Tsun-Hsuan Johnson Wang, Yen-Chen Lin, and Phillip Isola. 2022. Offline Multi-Agent Reinforcement Learning with Knowledge Distillation. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 226–237. https://proceedings.neurips.cc/paper_files/paper/2022/file/01d78b294d80491fecdde89cf03642-Paper-Conference.pdf
- [28] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *nature* 575, 7782 (2019), 350–354.
- [29] Shenzi Wang, Qisen Yang, Jiawei Gao, Matthieu Lin, HAO CHEN, Liwei Wu, Ning Jia, Shiji Song, and Gao Huang. 2023. Train Once, Get a Family: State-Adaptive Balances for Offline-to-Online Reinforcement Learning. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 47081–47104. https://proceedings.neurips.cc/paper_files/paper/2023/file/9318763d049edf9a1f2779b2a5991d3-Paper-Conference.pdf
- [30] Xihui Wang, Zheng Tian, Ziyu Wan, Ying Wen, Jun Wang, and Weinan Zhang. 2023. Order Matters: Agent-by-agent Policy Optimization. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=Q-neeWNv1>
- [31] Yiqin Yang, Xiaoteng Ma, Chenghao Li, Zewu Zheng, Qiyuan Zhang, Gao Huang, Jun Yang, and Qianchuan Zhao. 2021. Believe What You See: Implicit Constraint Approach for Offline Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 10299–10312. https://proceedings.neurips.cc/paper_files/paper/2021/file/550a141f12de6341fba65b0ad0433500-Paper.pdf
- [32] Haichao Zhang, Wei Xu, and Haonan Yu. 2023. Policy Expansion for Bridging Offline-to-Online Reinforcement Learning. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=Y34L45JR6z>
- [33] Yinmin Zhang, Jie Liu, Chuming Li, Yazhe Niu, Yaodong Yang, Yu Liu, and Wanli Ouyang. 2024. A Perspective of Q-value Estimation on Offline-to-Online Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [34] Kai Zhao, Jianye Hao, Yi Ma, Jinyi Liu, Yan Zheng, and Zhaopeng Meng. 2024. ENOTO: Improving Offline-to-Online Reinforcement Learning with Q-Ensembles. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, Kate Larson (Ed.). International Joint Conferences on Artificial Intelligence Organization, 5563–5571. <https://doi.org/10.24963/ijcai.2024/615>
- [35] Zhengbang Zhu, Minghuan Liu, Liyuan Mao, Bingyi Kang, Minkai Xu, Yong Yu, Stefano Ermon, and Weinan Zhang. 2023. MADiff: Offline Multi-agent Learning with Diffusion Models. *arXiv preprint arXiv:2305.17330* (2023).