

Novelty-Guided Data Reuse for Efficient and Diversified Multi-Agent Reinforcement Learning

Yangkun Chen*, Kai Yang*, Jian Tao, Jiafei Lyu†

Shenzhen International Graduate School, Tsinghua University
{chen-yk21, yk22, tj22, lvjf20}@mails.tsinghua.edu.cn

Abstract

Recently, deep Multi-Agent Reinforcement Learning (MARL) has demonstrated its potential to tackle complex cooperative tasks, pushing the boundaries of AI in collaborative environments. However, the efficiency of these systems is often compromised by inadequate sample utilization and a lack of diversity in learning strategies. To enhance MARL performance, we introduce a novel sample reuse approach that dynamically adjusts policy updates based on observation novelty. Specifically, we employ a Random Network Distillation (RND) network to gauge the novelty of each agent’s current state, assigning additional sample update opportunities based on the uniqueness of the data. We name our method Multi-Agent Novelty-GuidEd sample Reuse (MANGER). This method increases sample efficiency while promoting exploration and diverse agent behaviors. Our evaluations confirm substantial improvements in MARL effectiveness in complex cooperative scenarios such as Google Research Football and super-hard StarCraft II micromanagement tasks.

Code — https://github.com/kkane99/MANGER_code

Introduction

Reinforcement Learning (RL) (Sutton, Barto et al. 1998) has emerged as a powerful paradigm in artificial intelligence, enabling agents to learn optimal behaviors through interactions with their environment. In recent years, Multi-Agent Reinforcement Learning (MARL) (Canese et al. 2021) has attracted significant attention due to its applicability across various fields, including robotics (Wang et al. 2023; Guo et al. 2022; Duan et al. 2024), autonomous vehicles (Han et al. 2022; Zhang et al. 2022; Peng et al. 2021) and strategic games (Jia et al. 2020; Zhang et al. 2019; Vinyals et al. 2019; Yang et al. 2023).

MARL expands the principles of RL to scenarios involving multiple agents interacting with each other and the environment simultaneously. While traditional RL concentrates on a single agent that optimizes its behavior based on changes in the external environment, MARL necessitates

the simultaneous control of multiple agents. The algorithm must account for the complexity of coordinating behaviors among these agents, making the task more challenging in terms of difficulty, training duration, and convergence (Yang et al. 2020).

In MARL, each interaction with the environment incurs higher costs and time, making strategies more challenging to learn and converge. For most MARL algorithms, data is stored in a buffer following an interaction with the environment. A batch of data is then sampled from the buffer for updates before continuing with the environment interaction. Due to the task’s inherent complexity, immediately interacting with the environment after learning from a batch of data results in the strategy not fully utilizing the historical experience dataset. This leads to underfitting strategies during each interaction, resulting in no significant improvement in the quality of the sampled data and wasting interaction time and cost. To reduce the number of interactions with the environment and achieve desirable results, algorithms need to fully utilize historical data, train adequately, and then interact with the environment to obtain new, high-quality data.

Another critical aspect of MARL is the diversity among agents (Li et al. 2021; Bettini, Kortvelesy, and Prorok 2024). Unlike single-agent RL, which only needs to maximize its reward, some tasks require agents to play sacrificial roles, sacrificing individual gains for the best overall benefit. As seen in human societies, diverse skills, division of labor, and perspectives contribute to more effective collaboration. Similarly, in MARL, diverse agents can bring about a broader range of strategies and behaviors, ultimately leading to improved overall performance. Some recent work assigns specific roles to agents, executing certain actions to achieve diversity in agent strategies (Wang et al. 2020b; Hu et al. 2022). Others encourage exploration by maximizing the state entropy (Kim and Sung 2023; Tao et al. 2024) or other methods to diversify strategies among agents (Li et al. 2021; Yang et al. 2024a). However, these methods require additional improvements to the algorithm’s structure and do not perform well in complex environments.

To address these two key issues, we designed a method that improves the efficiency of sample utilization and improves diversity among agents. We have observed that the behavior exhibited by agents depends on the frequency of their policy updates. This phenomenon suggests that the up-

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: **Visualization of the environment.** In the SMAC image, the green box represents the tank role that actively absorbs damage and sacrifices itself to create an output environment for teammates. The red box represents the damage dealer role that activates attacks against enemies. The blue box represents the roaming role, similar to a guerrilla fighter, that can attract enemy aggro based on its own movement and lead some enemies away from the battlefield, preventing the enemies from focusing solely on the tank and causing it to be killed instantly.

date frequency of an agent’s policy plays a crucial role in shaping their behavior and subsequent performance in cooperative tasks. By simply controlling the update frequency of each agent, diverse strategies can be achieved among agents. Additionally, repeated updates allow for the reuse of samples, thereby improving sample utilization.

To calculate the number of updates for each agent, the update frequency for each agent is tailored based on the novelty of their current observations. Specifically, we use a measure of state diversity to gauge the novelty of observations encountered by each agent. Novel or rare states trigger more frequent policy updates for that agent, while common or familiar states result in fewer updates. This adaptive update schedule serves three main purposes. First, it enhances sample efficiency by allowing agents to reuse data multiple times, effectively extracting more information from the historically visited dataset. Second, we posit that the variations in the overall value function are primarily driven by the value functions of agents encountering novel states. Therefore, additional updates are necessary to these agents to ensure a more thorough fitting. Lastly, the diversity in the number of update steps for each agent allows agents to have different behavioral strategies, promoting diversity in strategies among different agents, and allowing each agent to learn specific skills and division of labor for better cooperation.

By improving sample utilization and promoting diverse behaviors among agents, our method aims to enhance their cooperative abilities and increase the success rate of tasks that require coordinated action. Through empirical evaluation and theoretical analysis, we have demonstrated the effectiveness of our method in improving the efficiency and performance of MARL systems across various domains.

Related Work

Multi-agent RL. The MARL field offers diverse approaches to tackle challenges like non-smoothness, efficient communication, and the balance between suboptimality and decentralization in multi-agent systems. A widely used architecture is Centralized Training with Decentralized Execution

(CTDE), where agents share historical data to enhance learning. CTDE assumes a central controller to process information from all agents, addressing the non-smoothness problem. Notable methods within this paradigm include MADDPG (Lowe et al. 2017), MAPPO (Yu et al. 2022), VDN (Sunehag et al. 2017), QMIX (Rashid et al. 2020), MAVEN (Mahajan et al. 2019), and QPLEX (Wang et al. 2020a). QMIX ensures monotonicity between global and individual Q-functions, while QPLEX relaxes the Individual-Global-Max (IGM) constraint with a duplex dueling network.

Research also focuses on extending MARL to handle high-dimensional, continuous state spaces via function approximators. Coordination graphs (Guestrin, Koller, and Parr 2001) factor large MDPs, enabling inter-agent communication through message passing, a concept supported by other studies (Lazaridou and Baroni 2020; Liu et al. 2020). To address the suboptimality-decentralization trade-off, methods like QTRAN (Son et al. 2019) introduce relaxed penalties, while QAtten (Yang et al. 2020) employs a multi-head attention-based Q-value mixing network. Finally, approaches like MAVEN (Mahajan et al. 2019) seek to encourage diverse behaviors through intrinsic rewards or hierarchical RL, maximizing mutual information between states and actions to foster agent diversity.

Sample Efficiency. In recent years, researchers have proposed various approaches to enhance sample efficiency in single or multi agent systems (Lyu et al. 2022, 2023, 2024a; Li et al. 2022; Yan, Lyu, and Li 2024; Gogineni et al. 2023). The Randomized Ensemble Double Q-learning (REDQ) algorithm (Chen, Hessel, and van Hasselt 2021) demonstrates that employing a larger Update-To-Data (UTD) ratio can significantly boost sample efficiency, yielding substantial performance improvements compared to model-based reinforcement learning algorithms. To mitigate potential overfitting issues, REDQ utilizes an ensemble of Q-networks to reduce estimation errors in Q-values. Similarly, the Adaptive Value-Targeted Learning (AVTD) algorithm (Oh, Hessel, and van Hasselt 2021) employs a validation set to estimate the fitting error of the current network, adopting a more

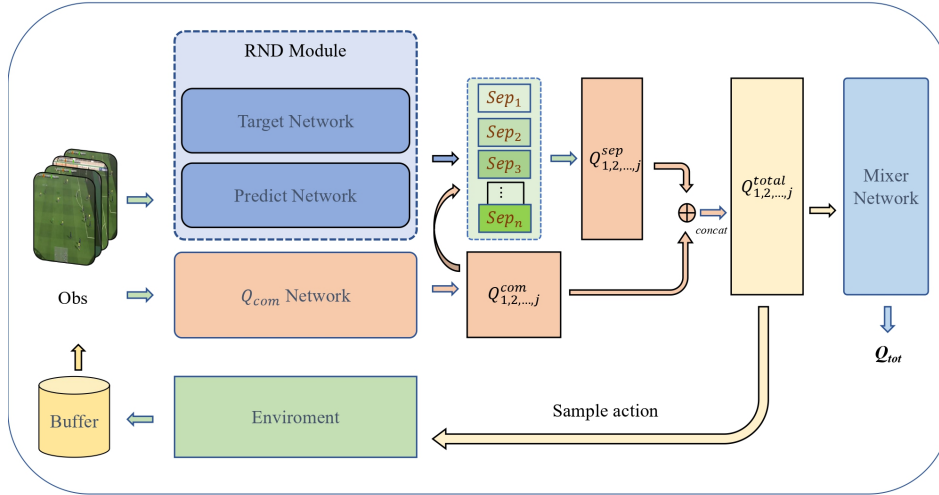


Figure 2: **Overview of the MANGER framework.** We employ the RND network to assess the novelty of each agent’s observations, thereby enabling differentiated updates among agents. Furthermore, we ensure that each additional update does not interfere with the agents by decomposing the network.

regularized approach with lower TD error to address overfitting problems associated with sample reuse. However, the Sample Multiple Reuse (SMR) algorithm (Lyu et al. 2024b) reveals that performing multiple repeated updates without an excessively large UTD value does not necessarily result in overfitting.

Agent Diversity. In MARL, agent diversity is crucial for enhancing system robustness, adaptability, and performance. This diversity, referring to variations in agents’ strategies, objectives, or learning processes, significantly influences multi-agent interactions. Recent research highlights the importance of fostering diversity to better address complex environments. One effective method to promote agent diversity is through role specialization, where agents adopt specialized behaviors to cover a wider range of strategies. Notable examples include RODE (Wang et al. 2020c), which dynamically identifies roles based on state and action history, allowing agents to adapt their strategies to the environment’s evolving needs. This method improves performance in cooperative tasks by facilitating scalable adaptation. ROMA (Wang et al. 2020b) also enhances diversity by using a latent variable model to infer roles and a mutual information objective to promote diversification, resulting in more effective cooperation. ASN (Wang et al. 2019) employs a behavior-based semantic neural network to calculate action semantics and achieve diversity by differentiating internal and external actions. The value of agent diversity is further emphasized in studies on emergent behaviors in unsupervised settings, where diverse behaviors can arise from simple reward structures in complex environments (Haber et al. 2018). This emergent diversity improves adaptability and system resilience. In competitive settings, training agents with diverse policies through population-based training leads to more robust strategies that generalize across var-

ious opponents (Baker et al. 2019). Additionally, DIYAN (Eysenbach et al. 2018) explores training agents to maximize behavioral diversity without explicit rewards, fostering a broad range of skills adaptable to various tasks.

In conclusion, agent diversity is essential for developing flexible and robust MARL systems. Our work aims to enhance adaptability and performance in dynamic environments by emphasizing the importance of diversity and effective discovery methods.

Preliminaries

Dec-POMDP: Our approach frames a fully cooperative multi-agent task within the context of a decentralized partially observable Markov decision process (Dec-POMDP) (Oliehoek, Amato et al. 2016), described by the tuple $M = \langle N, S, A, P, r, Z, O, \gamma \rangle$. In this formulation, N signifies a finite set of agents, $s_t \in S$ denotes the global state of the environment and $\gamma \in [0, 1)$ serves as the discount factor. At each time step, individual agents $j \in N$ perceive their own observations $o \in O$ and subsequently determine actions $a_k \in A$ based on the current global state s_t , where $k \in 1, 2, \dots, |A|$ delineates the action space’s capacity. These individual actions coalesce into a joint action vector a_t at time step t . The resultant joint reward, $r(s_t, a_t)$, triggers a transition in the environment as dictated by the transition function $P(s'|s, a_t)$. Subsequently, the joint policy π generates a joint action-value function: $Q_{\text{tot}}^\pi(s, a) = \mathbb{E}_{s_{t:\infty}, a_{t:\infty}} [G_t | s_t = s, a_t = a, \pi]$, where $G_t = \sum_{t'} \gamma^{t'} r_{t'+1}$ represents the expected discounted return. This comprehensive representation encapsulates the collaborative decision-making dynamics inherent in multi-agent systems.

Novelty We use the frequency of state visits as a measure of novelty. Many articles have analyzed this aspect of state novelty measurement, such as ICM (Pathak et al. 2017), RND

(Burda et al. 2018), CFN (Lobel, Bagaria, and Konidaris 2023), RCMP (Da Silva et al. 2020) and DRND (Yang et al. 2024b). In this paper, we employ the RND algorithm, a curiosity-driven mechanism, as the evaluation criterion for the novelty of current states or observations. Specifically, we have an untrained target network $f_{\text{target}}(o)$ and a network to be trained $f_{\text{predictor}}(o)$. During training, when the observation o has been seen many times, the mean squared error (MSE) loss between their predicted values tends to be small, and vice versa. Therefore, the higher the MSE loss, the higher the novelty of the observation. In this paper, we use this MSE loss to be the indicator of the novelty of each agent’s observation.

Method

In multi-agent reinforcement learning, almost all methods interact with the environment after updating parameters once with the data, without fully utilizing the samples. Furthermore, commonly used algorithms like QMIX, MAPPO, and MADDPG employ parameter-sharing techniques to reduce training costs and achieve some training effectiveness quickly, but they overlook the diversity of roles among agents, hindering the learning of more complex and better cooperative strategies. In this chapter, we propose MANGER based on the QMIX algorithm, which can determine the efficiency of sample utilization based on the novelty of each agent’s state, thereby allowing for different performances among different agents and improving sample utilization rates.

Assessing the Novelty of Observations

We employ RND to evaluate the novelty of agents’ observations. Initially, we initialize a trainable predictor network $f_{\text{predictor}}$ and a random, fixed target network f_{target} . When agents interact with the environment and receive observations, the novelty of observation o_i for agent i is calculated as follows:

$$N(o_i) = \|f_{\text{target}}(o_i) - f_{\text{predictor}}(o_i)\|^2. \quad (1)$$

The RND algorithm was originally designed to address the challenge of sparse rewards in environments by measuring the novelty of states and using intrinsic rewards to encourage agent exploration. However, this paper does not focus on exploration or sparse reward problems. RND is used solely as a metric to evaluate the novelty of observations.

Note that we utilize only a single total predictor and one target network instead of training individual predictors and preparing specialized target networks for each agent. This decision stems from the fact that when an observation o_i has been encountered multiple times, regardless of whether it was observed by other agents, we do not want to encourage the agent to consider it novel and reuse the data excessively. If an observation has been predominantly visited by agent i , reducing its novelty is reasonable as it indicates the observation has been visited frequently. Conversely, if the observation has been primarily visited by other agents, maintaining a low novelty prevents agent i from mimicking the behavior of those agents, thus enhancing the diversity among agents.

Subsequently, after computing the novelty of each observation across agents, we leverage this information to determine which data can be reused efficiently and how many times it is appropriate to reuse it.

Using Data Efficiently

Interacting with multi-agent environments is a time-consuming process, making the extra utilization of samples particularly important. From a holistic perspective, we enhance sample utilization rates by updating the overall framework with each data point twice. Additionally, considering each agent individually, we determine the novelty of data from each agent based on the aforementioned calculation of observation novelty, thereby deciding whether to reuse it. For data that has already appeared repeatedly, the network’s estimates are highly accurate, and thus further updating is unnecessary. Conversely, when agents encounter previously unseen observations, such as scoring a goal in a soccer environment or enemy units being destroyed in a StarCraft environment, since these states have rarely occurred before, we need to artificially update such data more frequently to ensure more accurate estimates under these observations. Additionally, leveraging these data more can encourage agents to explore these previously unseen states more, thereby enabling them to better learn challenging cooperative strategies and find globally optimal solutions. The criterion for performing additional updates is as follows: first, the novelty values $N_{\text{total}} = N_1, N_2, \dots, N_M$ within a batch are normalized, and then, based on the relationship between the normalized values and their variance, the number of additional updates is determined. The formula for the number of times samples are additionally updated is as follows:

$$T_i = \text{int} \left(\alpha \frac{(N_i - \bar{N}_{\text{total}})}{\text{Var}(N_{\text{total}})} \right). \quad (2)$$

Here, α is a coefficient controlling the rates of additional data reuse, \bar{N}_{total} represents the mean value of N_{total} , and $\text{Var}(N_{\text{total}})$ denotes the standard deviation of N_{total} . When T_i is less than 1, we do not perform extra updates on the network of agent i . In this study, we set $\alpha = 2$ and observe that the mean number of extra updates is less than 0.5, which does not significantly increase the training time.

Diversify the Agents by Separating the Critic Network

Due to the utilization of parameter-sharing techniques in most networks, such as QMIX, updating one agent triggers updates across all agent networks, which evidently does not align with our requirements. What we desire is diversity and specialization among each agent, so that even under the same observations, they exhibit distinct behaviors. However, if we allow parameters to differ among agents, both the parameter count and training difficulty would increase. To address this issue, inspired by the CDS algorithm(Li et al. 2021), we partition the critic network into shared and independent layers, and the Q value of agent i executing a_i in observation o_i is:

$$Q_{\text{tot}}^i(o_i, a_i) = Q_{\text{com}}(o_i, a_i) + \lambda Q_{\text{sep}}^i(o_i, a_i). \quad (3)$$

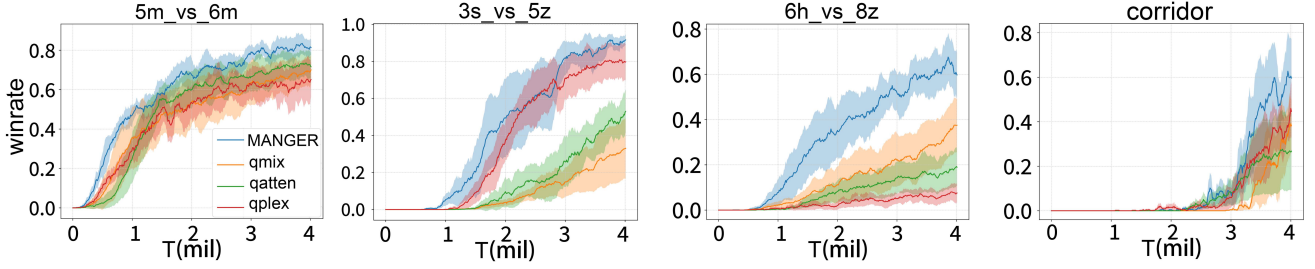


Figure 3: **Experimental results on SMAC.** All curves are averaged over 5 independent runs.




Figure 4: **Experiments on GRF environments.** All curves are averaged over 5 independent runs.

Here, λ controls the scale of separate Q values. The shared layer facilitates parameter sharing among all agents, aiding them in extracting overall features from the environment, which is similar to QMIX. This helps establish a consensus among agents regarding the value of observations and states, for instance, recognizing that attacking teammates is undesirable while attacking enemies is valuable. On the other hand, the independent layer embodies the diversity among agents. Taking a two-agent collaborative game as an example, some agents may play the role of “tank”, where executing aggressive actions in the current state is highly beneficial. Conversely, other agents playing the role of “attacker” may find aggressive actions unproductive, as there are already agents absorbing enemy fire, and thus opt not to attack. Hence, the independent layer serves to showcase individuality and roles among agents, with the final output of the critic network obtained through a weighted sum of the shared and independent layers.

In the process of performing additional training on agents, to ensure that only the current agent is updated without affecting others, we refrain from computing gradients for the shared layer during parameter updates and only update the network parameters of each agent’s independent layer. This approach not only avoids interdependence among agents but also speeds up the algorithm by performing partial gradient backpropagation. Only during overall data updates are both the shared and independent layers simultaneously updated across all agents.

Update of the MANGER Agents

The updated formulas for the total Q -value in QMIX and parameters are as follows:

$$Q_{\text{tot}}(s, a; \theta) = f_{\phi}(Q^1(o_1, a_1; \theta_1), \dots, Q^N(o_N, a_N; \theta_N)), \quad (4)$$

$$y = r(s, a) + \gamma \times \max_{a'} Q_{\text{tot}}(s', a'; \theta), \quad (5)$$

$$\theta \leftarrow \theta + \alpha \times (Q_{\text{tot}}(s, a; \theta) - y). \quad (6)$$

Here, f denotes the mixer network and α is the learning rate. By updating the parameters of the mixer network and the corresponding Q networks of each agent, the agents can estimate the overall Q value more accurately.

Comparing to QMIX, the update equations for MANGER agents are as follows:


$$Q_{\text{tot}}(s, a; \theta) = f_{\phi}(Q_{\text{tot}}^1(o_1, a_1; \theta_1), \dots, Q_{\text{tot}}^N(o_N, a_N; \theta_N)), \quad (7)$$

$$y = r(s, a) + \gamma \times \max_{a'} Q_{\text{tot}}(s', a'; \theta), \quad (8)$$


$$\theta \leftarrow \theta + \alpha \times (Q_{\text{tot}}(s, a; \theta) - y) \times H(T), \quad (9)$$

$$T \leftarrow T - 1. \quad (10)$$

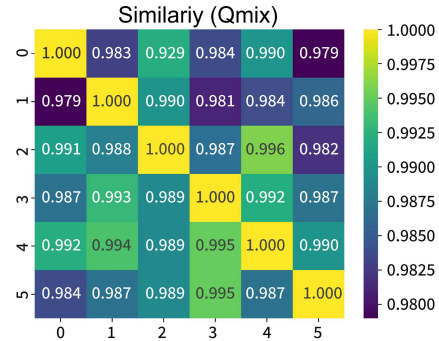
Here, $H(x) = [h(x_1), h(x_2), \dots, h(x_N)]$, $h(x) = \begin{cases} 1, & x > 0 \\ 0, & x \leq 0 \end{cases}$ and $T = [T_1, T_2, \dots, T_N]$ where T_i is calculated by equations (1) and (2). The formula described above will cyclically compute and continuously update the parameters until all components of T are not greater than 0. When



(a)



(b)



(c)

Figure 5: **Graphical illustration of agent diversity.** (a) shows how the agent within the red box should operate, with Q-values for actions. We calculated the Q-values of different agents and displayed three representative ones. For the same observation, Agent 1 moves southward to protect teammates; Agent 3 attacks low-health enemies; Agent 6 moves to the upper left to pull enemies away. We calculated the cosine similarity between the Q-values of other agents and the current agent’s Q-values. (b) and (c) compare our method with Qmix, showing cosine similarity with other agents under the same observation. Our agents exhibit more diversity.

the novelty corresponding to agent i is low, $h(T_i)$ will be 0 and will be ignored during parameter updates; when the degree of novelty is high, T_i will be greater than 1 and $h(T_i)$ will equal 1, which is equivalent to an additional update to the parameters θ_i corresponding to Q_i .

Experiment

In this section, we will experimentally address the following questions: (1) Can our method improve sample efficiency, leading to higher or faster convergence of agents’ win rates across different tasks? (2) Do agents exhibit diverse behaviors through updates of varying frequencies? (3) Is our proposed RND standalone update module effective compared to the holistic update? (4) Is our algorithm’s improvement due to adding the Q_{sep} module? To answer these four questions, we conducted the following experiments for validation.

Setup To validate the effectiveness of our proposed algorithm, we employed the widely used StarCraft Multi-Agent Challenge (SMAC, (Samvelyan et al. 2019)) in multi-agent reinforcement learning. SMAC includes various environments that require cooperation and coordination and is based on the popular real-time strategy game StarCraft II, providing a diverse set of tasks for multi-agent collaboration. We also used the Google Research Football (GRF) (Kurach et al. 2020) environment, which contains numerous multi-

agent tasks where agents must cooperate to achieve goals, serving as a testbed for assessing the effectiveness of algorithms in handling complex, real-world-like multi-agent scenarios. We utilized PyMARL2 (Hu et al. 2021) as our codebase and employed QMIX as our baseline algorithm, evaluating whether QMIX combined with our method yields superior performance. We compared our approach against several popular methods, including QMIX, QPLEX, and Qatten, using the parameters recommended in the respective papers.

Result In this section, we will address the question (1) and validate whether our algorithm can enhance the convergence rate of agents, thus accelerating the improvement of win rates. Within the SMAC environment, we have chosen a variety of tasks to evaluate the effectiveness of our algorithm, including both symmetric and asymmetric environments. We recorded the test win rates of each method on various tasks and compared the final performance and convergence rates of different methods. We plotted win rate curves of different methods under various task environments for comparison, as shown in Figure 3. It can be observed that in the majority of environments, our method is able to significantly enhance the performance of the baseline QMIX algorithm within a short period. In the GRF environment, we selected three more challenging settings, and the experimental results are shown in Figure 2. It can be seen that our

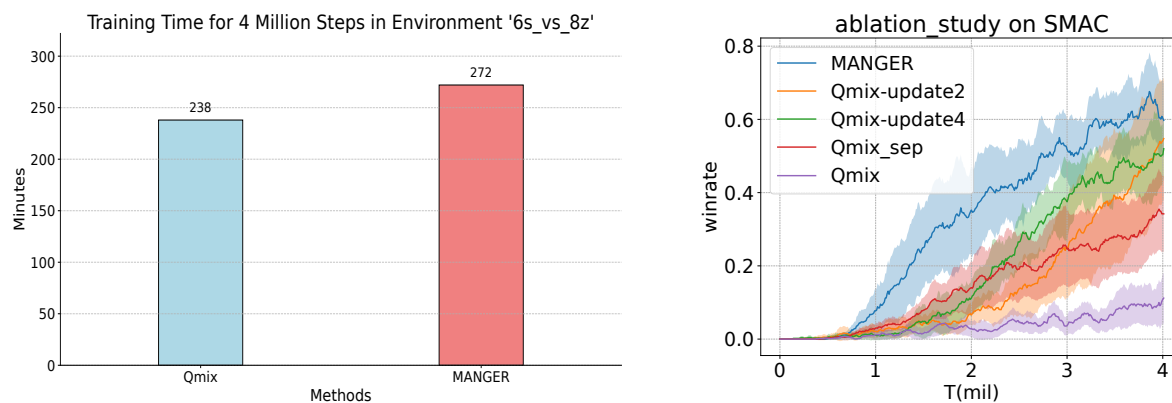


Figure 6: **Left:** Training time comparison between the MANGER and Qmix methods in the 6h_vs_8z environment. **Right:** Ablation study of MANGER in the 6h_vs_8z environment. It can be observed that our method does not significantly increase training time while achieving performance improvements with the same number of environment interactions. The ablation study further confirms that the performance gains are not solely due to increased sample reuse or network decomposition but are specifically attributed to the targeted utilization of certain important samples by the RND module.

method performs the best among all the methods. By leveraging sample reuse, our method enables agents to better utilize data, thereby improving task win rates, thus validating the effectiveness of our algorithm. Furthermore, compared to other updated improvement algorithms like Qatten and QPLEX, our method achieves the best performance across all methods, further demonstrating its effectiveness.

Analysis To answer the question (2), whether agents have demonstrated differentiated performance, we conducted both qualitative and quantitative analyses, as shown in Figure 5. During the visualization evaluation, we selected a representative frame to observe the Q values of different agents under the same observation and their corresponding strategies. It can be seen that under the same observation, the Q values of different agents trained by MANGER have significant differences, and the actions they take are also different. For the QMIX algorithm, this would be impossible, as the Q values would be very close under the same observation. Through a quantitative comparison using heat maps, it can be seen that the Q value similarities among different agents trained by our method are lower, showing more differentiation and division of labor, which proves the effectiveness of our algorithm. In addition, we have also visualized the environment to observe whether our algorithm can enable agents to achieve division of labor and generate diversity, as shown in Figure 1. It can be seen that agents can take on the role of a tank that actively absorbs damage to help teammates with output; activate the offensive role of attackers; attract hatred to pull monsters, avoiding the continuous activation of the tank by roaming characters.

Ablation Study In this section, we will verify questions (3) and (4). To validate the effectiveness of individually updating each datum with the RND module, we conducted statistical analysis on the average individual update counts. We found that due to the nature of the normal distribution, the extra average update times tend to be about 0.5, which does not significantly increase the training cost. We conducted

comparisons in some environments, and the comparison of training time and the final results are shown in Figure 6. It can be observed that compared to QMIX, QMIX with an added separate value, and methods that fix multiple updates on this basis, our method achieves better training results with relatively fewer updates. This demonstrates that utilizing the RND network for individual updates enhances effectiveness not merely by increasing the average additional update counts, but rather by targeting additional updates towards crucial states and trajectories. Additionally, the performance of adding a separate module to QMIX not being as good as MANGER also proves that the previous results were not due to the introduction of new modules, but rather the outcome of repeated updates by the RND module.

Conclusion

To improve sample efficiency in multi-agent systems and enhance performance diversity among different agents, enabling them to develop specialization for collaborative task completion, this paper proposes the MANGER method to address this issue. By using the RND method to calculate the novelty of each agent’s observations, states with high novelty and inaccurate Q-value estimates are updated multiple times, while states with low novelty are not given additional updates. This approach enables differentiated updates for different agents, increases sample efficiency, and allows for targeted extra updates on rare samples. Experimental results show that our method, which only provides additional updates to novel states, does not introduce significant additional time overhead. In SMAC and GRF environments, our method outperforms QMIX, Qatten, and Qplex algorithms, achieving higher win rates within fewer training steps. By observing the similarity of Q-values among different agents and through visualization, we demonstrate that MANGER can enhance the diversity of agent policies and enable specialization for collaborative task completion, proving the effectiveness and feasibility of our method.

Acknowledgements

This work was supported by the STI 2030-Major Projects under Grant 2021ZD0201404. The authors also thank the anonymous reviewers for valuable comments.

References

- Baker, B.; Kanitscheider, I.; Markov, T.; Wu, Y.; Powell, G.; McGrew, B.; and Mordatch, I. 2019. Emergent tool use from multi-agent autotutorials. *arXiv preprint arXiv:1909.07528*.
- Bettini, M.; Kortvelesy, R.; and Prorok, A. 2024. Controlling Behavioral Diversity in Multi-Agent Reinforcement Learning. *ArXiv*, abs/2405.15054.
- Burda, Y.; Edwards, H.; Storkey, A.; and Klimov, O. 2018. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*.
- Canese, L.; Cardarilli, G. C.; Nunzio, L. D.; Fazzolari, R.; Giardino, D.; Re, M.; and Spanò, S. 2021. Multi-Agent Reinforcement Learning: A Review of Challenges and Applications. *Applied Sciences*, 11: 4948.
- Chen, M.; Hessel, M.; and van Hasselt, H. 2021. Randomized Ensemble Double Q-learning: Learning Fast Without a Model. *Advances in Neural Information Processing Systems*, 34: 14527–14539.
- Da Silva, F. L.; Hernandez-Leal, P.; Kartal, B.; and Taylor, M. E. 2020. Uncertainty-aware action advising for deep reinforcement learning agents. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 5792–5799.
- Duan, K.; Yang, K.; Liu, H.; and Wang, X. 2024. Novelty-based Sample Reuse for Continuous Robotics Control. *arXiv preprint arXiv:2410.13490*.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*.
- Gogineni, K.; Mei, Y.; Lan, T.; Wei, P.; and Venkataramani, G. 2023. Accmer: Accelerating multi-agent experience replay with cache locality-aware prioritization. In *2023 IEEE 34th International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, 205–212. IEEE.
- Guestrin, C.; Koller, D.; and Parr, R. 2001. Multiagent planning with factored MDPs. *Advances in neural information processing systems*, 14.
- Guo, Y. S.; Campbell, J.; Stepputtis, S.; Li, R.; Hughes, D.; Fang, F.; and Sycara, K. P. 2022. Explainable Action Advising for Multi-Agent Reinforcement Learning. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 5515–5521.
- Haber, N.; Mrowca, D.; Wang, S.; Fei-Fei, L. F.; and Yamins, D. L. 2018. Learning to play with intrinsically-motivated, self-aware agents. *Advances in neural information processing systems*, 31.
- Han, S.; Wang, H.; Su, S.; Shi, Y.; and Miao, F. 2022. Stable and Efficient Shapley Value-Based Reward Reallocation for Multi-Agent Reinforcement Learning of Autonomous Vehicles. *2022 International Conference on Robotics and Automation (ICRA)*, 8765–8771.
- Hu, J.; Jiang, S.; Harding, S. A.; Wu, H.; and wei Liao, S. 2021. Rethinking the Implementation Tricks and Monotonicity Constraint in Cooperative Multi-Agent Reinforcement Learning.
- Hu, S.; Xie, C.; Liang, X.; and Chang, X. 2022. Policy Diagnosis via Measuring Role Diversity in Cooperative Multi-agent RL. *ArXiv*, abs/2207.05683.
- Jia, H.; Hu, Y.; Chen, Y.; Ren, C.; Lv, T.; Fan, C.; and Zhang, C. 2020. Fever Basketball: A Complex, Flexible, and Asynchronous Sports Game Environment for Multi-agent Reinforcement Learning. *ArXiv*, abs/2012.03204.
- Kim, W.; and Sung, Y. 2023. An adaptive entropy-regularization framework for multi-agent reinforcement learning. In *International Conference on Machine Learning*, 16829–16852. PMLR.
- Kurach, K.; Raichuk, A.; Stańczyk, P.; Zajac, M.; Bachem, O.; Espeholt, L.; Riquelme, C.; Vincent, D.; Michalski, M.; Bousquet, O.; et al. 2020. Google Research Football: A Novel Reinforcement Learning Environment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 4501–4510.
- Lazaridou, A.; and Baroni, M. 2020. Emergent multi-agent communication in the deep learning era. *arXiv preprint arXiv:2006.02419*.
- Li, C.; Wang, T.; Wu, C.; Zhao, Q.; Yang, J.; and Zhang, C. 2021. Celebrating diversity in shared multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 3991–4002.
- Li, X.; Qiao, Z.; Gong, A.; Lyu, J.; Yu, C.; Yan, J.; and Li, X. 2022. Prag: Periodic regularized action gradient for efficient continuous control. In *Pacific Rim International Conference on Artificial Intelligence*, 106–119. Springer.
- Liu, Y.-C.; Tian, J.; Glaser, N.; and Kira, Z. 2020. When2com: Multi-agent perception via communication graph grouping. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 4106–4115.
- Lobel, S.; Bagaria, A.; and Konidaris, G. 2023. Flipping Coins to Estimate Pseudocounts for Exploration in Reinforcement Learning. *arXiv preprint arXiv:2306.03186*.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Lyu, J.; Bai, C.; Yang, J.; Lu, Z.; and Li, X. 2024a. Cross-domain policy adaptation by capturing representation mismatch. *arXiv preprint arXiv:2405.15369*.
- Lyu, J.; Ma, X.; Yan, J.; and Li, X. 2022. Efficient continuous control with double actors and regularized critics. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 7655–7663.
- Lyu, J.; Wan, L.; Li, X.; and Lu, Z. 2024b. Off-policy RL algorithms can be sample-efficient for continuous control via sample multiple reuse. *Information Sciences*, 666: 120371.
- Lyu, J.; Yang, Y.; Yan, J.; and Li, X. 2023. Value activation for bias alleviation: Generalized-activated deep double deterministic policy gradients. *Neurocomputing*, 518: 70–81.

- Mahajan, A.; Rashid, T.; Samvelyan, M.; and Whiteson, S. 2019. Maven: Multi-agent variational exploration. *Advances in neural information processing systems*, 32.
- Oh, J.; Hessel, M.; and van Hasselt, H. 2021. Adaptive Value-Targeted Learning. *Proceedings of the 38th International Conference on Machine Learning (ICML)*.
- Oliehoek, F. A.; Amato, C.; et al. 2016. *A concise introduction to decentralized POMDPs*, volume 1. Springer.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, 2778–2787. PMLR.
- Peng, Z.; Li, Q.; Hui, K. M.; Liu, C.; and Zhou, B. 2021. Learning to simulate self-driven particles system with coordinated policy optimization. *Advances in Neural Information Processing Systems*, 34: 10784–10797.
- Rashid, T.; Samvelyan, M.; De Witt, C. S.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *The Journal of Machine Learning Research*, 21(1): 7234–7284.
- Samvelyan, M.; Rashid, T.; De Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*.
- Son, K.; Kim, D.; Kang, W. J.; Hostallero, D. E.; and Yi, Y. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International conference on machine learning*, 5887–5896. PMLR.
- Sunehag, P.; Lever, G.; Grusl, A.; Czarnecki, W. M.; Zambaldi, V.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; et al. 2017. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296*.
- Sutton, R. S.; Barto, A. G.; et al. 1998. Introduction to reinforcement learning.
- Tao, J.; Chen, Y.; Zhang, Y.; Yang, K.; and Li, X. 2024. Multi-agent Exploration with Sub-state Entropy Estimation. In *2024 International Joint Conference on Neural Networks (IJCNN)*, 1–9. IEEE.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D.; Powell, R.; Ewalds, T.; Georgiev, P.; Oh, J.; Horgan, D.; Kroiss, M.; Danihelka, I.; Huang, A.; Sifre, L.; Cai, T.; Agapiou, J. P.; Jaderberg, M.; Vezhnevets, A. S.; Leblond, R.; Pohlen, T.; Dalibard, V.; Budden, D.; Sulsky, Y.; Molloy, J.; Paine, T. L.; Gulcehre, C.; Wang, Z.; Pfaff, T.; Wu, Y.; Ring, R.; Yogatama, D.; Wünsch, D.; McKinney, K.; Smith, O.; Schaul, T.; Lillicrap, T. P.; Kavukcuoglu, K.; Hassabis, D.; Apps, C.; and Silver, D. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575: 350 – 354.
- Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2020a. Qplex: Duplex dueling multi-agent q-learning. *arXiv preprint arXiv:2008.01062*.
- Wang, T.; Dong, H.; Lesser, V.; and Zhang, C. 2020b. Roma: Multi-agent reinforcement learning with emergent roles. *arXiv preprint arXiv:2003.08039*.
- Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2020c. Rode: Learning roles to decompose multi-agent tasks. *arXiv preprint arXiv:2010.01523*.
- Wang, W.; Mao, L.; Wang, R.; and Min, B.-C. 2023. Multi-Robot Cooperative Socially-Aware Navigation Using Multi-Agent Reinforcement Learning. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 12353–12360.
- Wang, W.; Yang, T.; Liu, Y.; Hao, J.; Hao, X.; Hu, Y.; Chen, Y.; Fan, C.; and Gao, Y. 2019. Action semantics network: Considering the effects of actions in multiagent systems. *arXiv preprint arXiv:1907.11461*.
- Yan, M.; Lyu, J.; and Li, X. 2024. Enhancing visual reinforcement learning with State-Action Representation. *Knowledge-Based Systems*, 304: 112487.
- Yang, K.; Fang, Z.; Li, X.; and Tao, J. 2024a. Cmbe: Curiosity-driven model-based exploration for multi-agent reinforcement learning in sparse reward settings. In *2024 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Yang, K.; Gong, A.; Tao, J.; Zhang, Y.; and Li, X. 2023. GTLMA: Generalizable Hierarchical Learning for Tasks with Variable Entities. In *2023 International Conference on Frontiers of Robotics and Software Engineering (FRSE)*, 97–103. IEEE.
- Yang, K.; Tao, J.; Lyu, J.; and Li, X. 2024b. Exploration and Anti-Exploration with Distributional Random Network Distillation. *arXiv preprint arXiv:2401.09750*.
- Yang, Y.; Hao, J.; Liao, B.; Shao, K.; Chen, G.; Liu, W.; and Tang, H. 2020. Qatten: A general framework for cooperative multiagent reinforcement learning. *arXiv preprint arXiv:2002.03939*.
- Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35: 24611–24624.
- Zhang, Z.; Han, S.; Wang, J.; and Miao, F. 2022. Spatial-Temporal-Aware Safe Multi-Agent Reinforcement Learning of Connected Autonomous Vehicles in Challenging Scenarios. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 5574–5580.
- Zhang, Z.; Li, H.; Zhang, L.; Zheng, T.; Zhang, T.; Hao, X.; Chen, X.; Chen, M.; Xiao, F.; and Zhou, W. 2019. Hierarchical Reinforcement Learning for Multi-agent MOBA Game. *ArXiv*, abs/1901.08004.