

Self-Motivated Multi-Agent Exploration

Shaowei Zhang^{*}

National Key Laboratory for Novel
Software Technology, Nanjing
University
Nanjing, P.R.China
zhangsw@lamda.nju.edu.cn

Jiahao Cao^{*}

National Key Laboratory for Novel
Software Technology, Nanjing
University
Nanjing, P.R.China
caojh@lamda.nju.edu.cn

Lei Yuan

National Key Laboratory for Novel
Software Technology, Nanjing
University; Polixir Technologies
Nanjing, P.R.China
yuanl@lamda.nju.edu.cn

Yang Yu

National Key Laboratory for Novel
Software Technology, Nanjing
University; Polixir Technologies
Nanjing, P.R.China
yuy@nju.edu.cn

De-Chuan Zhan[†]

National Key Laboratory for Novel
Software Technology, Nanjing
University; Polixir Technologies
Nanjing, P.R.China
zhandc@nju.edu.cn

ABSTRACT

In cooperative multi-agent reinforcement learning (CMARL), it is critical for agents to achieve a balance between self-exploration and team collaboration. However, agents can hardly accomplish the team task without coordination and they would be trapped in a local optimum where easy cooperation is accessed without enough individual exploration. Recent works mainly concentrate on agents' coordinated exploration, which brings about the exponentially grown exploration of the state space. To address this issue, we propose *Self-Motivated Multi-Agent Exploration* (SMMAE), which aims to achieve success in team tasks by adaptively finding a trade-off between self-exploration and team cooperation. In SMMAE, we train an independent exploration policy for each agent to maximize their own visited state space. Each agent learns an adjustable exploration probability based on the stability of the joint team policy. The experiments on highly cooperative tasks in StarCraft II micromanagement benchmark (SMAC) demonstrate that SMMAE can explore task-related states more efficiently, accomplish coordinated behaviours and boost the learning performance.

KEYWORDS

Multi-agent Reinforcement Learning; Self-motivated Exploration; Multi-agent Cooperation

ACM Reference Format:

Shaowei Zhang, Jiahao Cao, Lei Yuan, Yang Yu, and De-Chuan Zhan. 2023. Self-Motivated Multi-Agent Exploration. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 11 pages.

1 INTRODUCTION

Cooperative *multi-agent reinforcement learning* (CMARL) has achieved outstanding results [13, 25] and has been applied to many real-world applications, such as autonomous driving [54], multi-agent path

^{*} Equal contribution.

[†] Corresponding author.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

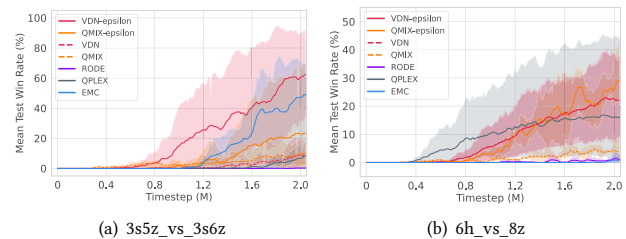


Figure 1: Results on two super hard maps need strong exploration from SMAC, VDN-epsilon and QMIX-epsilon can promote coordination significantly.

finding [12], natural language processing tasks [42], and dynamic algorithm configuration [47]. Most of the methods can be divided into two categories: the value-based methods [34, 36, 43] and others based on policy gradient [10, 20, 49]. The value-based methods and their variants [6, 50, 51] can have better performance in complex tasks like the StarCraft II micromanagement benchmark [35]. However, they usually neglect efficient exploration, which is a particularly challenging problem in complex scenarios. Agents will be trapped in a local optimum if they only focus on team collaboration and cannot conduct coordinated behaviours if they only pay attention to exploration.

To address the problem of MARL exploration, a vanilla technique that is commonly used in value-based MARL algorithms is ϵ -greedy [13, 48]. Recently, a variety of methods have been proposed and they focus more on coordinated exploration. EITI & EDTI [45] use the interactions between the agents to measure the collaborative exploration. CMAE [19] focuses on the low-dimensional space that indeed affects the team cooperation reward. EMC [53] proposes that the local Q function is affected by the influence of other agents and guides the exploration based on this viewpoint. However, these algorithms ignore individual exploration and consider only joint exploration objective, which has a very large space and is not as efficient as individual exploration.

The mentioned methods can facilitate the MARL exploration ability in some way but are inefficient in complex scenarios, as those methods seldom consider self-exploration from an individual point of view. What's more, they are far away from human coordination.

If a person is not familiar with the team behaviour in the current situation, one will first learn how to cooperate with others. Otherwise, one will explore more unknown situations to cooperate better with others in more situations [16]. Although those mentioned MARL exploration approaches can achieve coordination improvements in some MARL tasks by designing complex modules, we find that we can achieve competitive results with vanilla MARL methods such as QMIX [34], by simply focusing on self-exploration.

For example, in the widely-used benchmark SMAC [35] almost all existing MARL methods fail to succeed in 2 million steps in the challenging scenarios such as maps 3s5z_vs_3s6z and 6h_vs_8z, where sufficient exploration is required [34, 36, 43, 44, 53]. However, in our experiments, by adjusting the degree of self-exploration, we find that QMIX [34] and VDN [36] can achieve significant improvements in exploration and can succeed in super hard tasks (Figure 1). Concretely, here the self-exploration is adjusted by changing the ε end value of ε -greedy (Figure 12 in Appendix C) which starts with the value 1.0 and decays linearly with time to a fixed end value ε_T in the popular implementation of QMIX [35]. The results show that the ε value has a strong association with the degree of exploration, and the probability that the agent adopts an exploration policy at one step during training time, should be applied at different timesteps during training to trade off between self-exploration and team cooperation to for high coordination. However, to achieve this, it needs extensive work for hyper-parameter tuning at each timestep, and it is not tractable to manually decide which ε value is suitable at each timestep in complex tasks.

Towards designing a method that can adaptively adjust the degree of exploration at different timesteps, we propose a novel multi-agent exploration method *Self-Motivated Multi-Agent Exploration* (SMMAE). Inspired by the widely used uncertainty measurement from single agent Reinforcement Learning [18, 31], we posit that learning a proper uncertainty about the multi-agent system and employing it to promote individual exploration can facilitate coordination in CMARL. In a multi-agent system, when the uncertainty between agents' actions and others' observations is limited, the agents should explore individually to jump out of the local optimum. On the contrary, the agents' lack of awareness of others makes it hard to achieve coordinated behaviors when the uncertainty in the multi-agent system is high, and the agents should explore less and learn how to cooperate first before exploring more. Specifically, the uncertainty of the joint policy for the agents can be measured using the correlation between each agent's action and observations of other agents. We then use mutual information [17] to denote the correlation. The observation of other agents is used to predict the current agent's action, and the cross entropy loss is used as a criterion. The smaller the cross entropy loss is, the less the uncertainty is. Our main contributions are:

- We study the exploration probability in MARL and reveal that the suitable ε value at different timestep has a huge valuable influence on exploration behaviours and final performance, especially in complex environments.
- SMMAE can adaptively adjust exploration probability according to the uncertainty in the multi-agent system to trade off between self-exploration and team cooperation.
- The experiments demonstrate the strong ability of SMMAE to explore task-related state space efficiently.

2 BACKGROUND

2.1 Cooperative Multi-Agent System Model

This paper considers a fully cooperative multi-agent system, where all the agents need to cooperative with each other to earn a shared team reward. The system can be modelled by a Dec-POMDP [23] tuple $G = \langle \mathcal{N}, \mathcal{S}, \mathcal{A}, P, \mathcal{R}, \Omega, O, n, \gamma \rangle$, where \mathcal{N} is a finite set of n agents, $s \in \mathcal{S}$ is a global state of the environment in the set of possible states, \mathcal{A} is the finite action set and $\gamma \in [0, 1)$ is the discount factor. Due to the partially observable settings, agent $i \in \mathcal{N}$ is only accessible to a local observation $o_i \in \Omega$ according to the observation function $O(s, i)$. Each agent has an observation-action trajectory history $\tau_i \in \mathcal{T} \equiv (\Omega \times \mathcal{A})^*$. At each timestep the joint action of the team is $\mathbf{a} = \langle a_1, \dots, a_n \rangle \in \mathcal{A} \equiv \mathcal{A}^n$ where $a_i \in \pi_i(a | \tau_i)$ is an action selected by each agent i . The team joint action will lead to the next global state s' according to the environment transition function $P(s' | s, \mathbf{a})$ and the team will earn the shared reward $r = \mathcal{R}(s, \mathbf{a})$. The joint policy $\pi \equiv \langle \pi_1, \dots, \pi_n \rangle$ aims to maximize the joint value function $V^\pi(s) = \mathbb{E}_{s_0: \infty} [\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, \pi]$, which induces a joint action-value function $Q^\pi(s, \mathbf{a}) = \mathcal{R}(s, \mathbf{a}) + \gamma \mathbb{E}_{s'} [V^\pi(s')]$.

2.2 Centralized Training with Decentralized Execution (CTDE) paradigm

Centralized Training with Decentralized Execution (CTDE) [24] has been a popular paradigm for cooperative multi-agent reinforcement learning. In a CTDE paradigm, during training the global states are available for all the agents by using a centralized controller, while during test time each agent has to use its local network to select an action based on its local trajectories. Deep Q Network [22] and its derivatives [40, 46] have achieved great performance in reinforcement learning tasks, especially in game tasks. In the multi-agent system, the tuple in replay buffer \mathcal{D} is $(\tau_t, \mathbf{a}_t, r_t, \tau_{t+1})$ and the joint action-value function is $Q_{tot}(\tau_t, \mathbf{a}_t; \theta)$, where θ are the parameters of the Q network and will be learnt by the following Temporal Difference (TD) error [37]:

$$\begin{aligned} \mathcal{L}(\theta) &= \mathbb{E}_{(\tau_t, \mathbf{a}_t, r_t, \tau_{t+1}) \sim \mathcal{D}} \left[r + \gamma \max_{\mathbf{a}_{t+1}} Q_{tot}(\tau_{t+1}, \mathbf{a}_{t+1}; \theta^-) - Q_{tot}(\tau_t, \mathbf{a}_t; \theta) \right]^2, \end{aligned} \quad (1)$$

where θ^- are the parameters of the target network and will be updated by θ periodically. Considering the CTDE paradigm, many works [33, 34, 36, 43] adopt the decomposition structure between the joint action-value function Q_{tot} and local action-value function Q_i , and obtain good performance.

3 RELATED WORK

3.1 Exploration in Reinforcement Learning

Various methods for exploration have been studied in single-agent reinforcement learning [48]. Some works explore by focusing on the environment dynamics. ICM [28] adopts both forward and inverse models to build a good feature space for curiosity and explore what is controlled by or affects the agent. Pathak et al. [29] propose to use an ensemble of dynamics models. They use the variance over the output of these networks in the ensemble to induce the exploration. There are some works focused on the environment novelty to induce exploration. Some works are the generation of

count-based methods and use density models which can generate pseudo-counts of visited states to measure the uncertainty of the agents [3, 26]. Some works use prediction error to reflect the novelty of the states. RND [5] uses a fixed randomly initialized network to get the embedding of the state, then uses a learnable network to reconstruct this embedding, and the reconstruction loss will be adopted as an intrinsic reward to guide exploration. Based on this, NovelD [52] proposes to use the difference of RND novelty between adjacent time step as the intrinsic reward, and adds a restriction in an episode so that the agent could earn a reward only when it visits the states for the first time. Therefore, NovelD prefers to explore unexplored boundary states, thereby exploring more valuable state regions. SMM [18] aims to learn a state marginal distribution that matches the prior distribution and uses this as the target of exploration. Some works adopt epsilon schedule for better exploration. Tokic [38] uses the difference between value functions to adjust the size of epsilon. When the agent’s knowledge becomes certain about the environment, the degree of exploration will be reduced. ϵ -z-greedy [8] replaces a single action with a sequence of actions (option), choosing an option instead of choosing a random action with probability ϵ . Some works focus on noise for better exploration. NoisyNet [11] replaces the fully connected value network with a learnable noise network. Plappert et al. [32] add adaptive-scale noise to the parameters of the policy. There are also some works that study exploration from some interesting perspectives. Eysenbach et al. [9] use diversity-driven exploration to learn distinguishable and diverse skills, while SMiRL [4] reduces the entropy of the visited states during exploration to exclude negative effects of perturbations and acquire complex behaviours and skills without supervision. C-BET [27] first learns a exploration policy across environments without extrinsic rewards, then transfers the learned exploration policy to the target tasks. Pislár et al. [30] study the problem when the agent should explore and reveal the results of the methods with different types of exploring temporal structure. Agent57 [1] adopts a more engineering approach. The Q network is divided into an intrinsic part and an extrinsic part. NGU is a policy that treats different degrees of exploration equally [2]. Based on NGU, Agent57 uses a meta-controller to select the policy adaptively.

3.2 Multi-Agent Exploration

There are also many studies on exploration in multi-agent reinforcement learning. To achieve committed exploration, MAVEN [21] adopts hierarchical control and the policies of the agents are conditioned on the shared latent variable generated by a hierarchical policy. Wang et al. [45] propose two exploration methods, EITI and EDTI, to induce cooperative exploration by capturing the influence of one agent’s on other agents. EITI quantifies the influences on the state transition dynamics, while EDTI quantifies both transition and reward influences. However, they are not scalable as they need to use a approximation way to measure the influence of other agents on one agent when there are many agents, which can cause the method to fail. CMAE [19] proposes that reward function only depends on a small subset of the large state space. Therefore, it first explores in the projected low-dimensional space of the high-dimensional state space, then they select goals from the low-dimensional space and train the exploration policies to reach the goal to explore higher-dimensional space continuously.

However, it’s hard to find the effective low-dimensional projection in complex MARL tasks. EMC [53] proposes that local Q function of each agent can capture the novelty of states and the influences between agents. To induce coordinated exploration, EMC proposes to use prediction errors of individual Q function as the intrinsic reward. However, EMC is also not scalable, as it has to maintain a huge episodic memory buffer.

4 METHODOLOGY

In this section, we introduce *Self-Motivated Multi-Agent Exploration* (SMMAE), a novel method for effective exploration in MARL. Figure 2 shows the whole SMMAE framework. Part (a) and (b) in Figure 2 are the common structures in Centralized Training with Decentralized Execution (CTDE) algorithms using value decomposition. Each agent i uses its observation o_t^i and action a_{t-1}^i as input and outputs local Q-value to select action. The mixing network takes all the outputs of the agents to train the multi-agent policy during training time. SMMAE mainly changes the structure of each individual agent in part (b), and the main contribution of our method is illustrated in part (c).

4.1 Adaptive Exploration

In ϵ -greedy, the value of ϵ will decay linearly to a fixed end value ϵ_T in a common implementation. However, the experiments prove that exploration probability with fixed end value is not efficient (Figure 1). Therefore, we adopt adaptive exploration probability. As the exploration probability ϵ is for all agents, we use $\alpha = \langle \alpha_1, \dots, \alpha_n \rangle$ to denote the exploration probabilities of n agents, where α_i represents the probability of agent i to select the exploration policy.

We associate α with the uncertainty of the multi-agent system (Part (c) in Figure 2). The agents cooperate well when the uncertainty of the multi-agent system is limited and they should increase the exploration intensity to jump out of the local optimum by increasing α_i . When the system uncertainty is high, the agents lack awareness of others and are hard to achieve coordinated behaviours. Under this condition, agents should take more exploitation to reduce the uncertainty by decreasing α_i .

The uncertainty of the multi-agent system can be measured by the correlation between agents. When the correlation between agent i ’s action a_i and the other agents’ observation o_{-i} is low, the uncertainty of the multi-agent system is high because the agents focus more on their own action selection than the team cooperation. We propose to use the mutual information $\mathcal{I}(a_i, o_{-i})$ to reflect the correlation between the agent i ’s action a_i and the other agents’ observation o_{-i} :

$$F(\mathbf{o}, \mathbf{a}) = \sum_{i=1}^n \mathcal{I}(a_i, o_{-i}) = \sum_{i=1}^n \mathcal{H}(a_i) - \mathcal{H}(a_i | o_{-i}), \quad (2)$$

where \mathcal{H} denotes entropy. Then we can derive a lower bound for the team mutual information term using a variational estimator:

$$\begin{aligned} F(\mathbf{o}, \mathbf{a}) &= \sum_{i=1}^n \mathbb{E}_{a_i \sim p(a_i | o_{-i})} [\log p(a_i | o_{-i})] - \mathbb{E}_{a_i \sim p(a_i)} [\log p(a_i)] \\ &\geq \sum_{i=1}^n \mathbb{E}_{a_i \sim p(a_i | o_{-i})} [\log q_{\xi_i}(a_i | o_{-i})] - \mathbb{E}_{a_i \sim p(a_i)} [\log p(a_i)], \end{aligned} \quad (3)$$

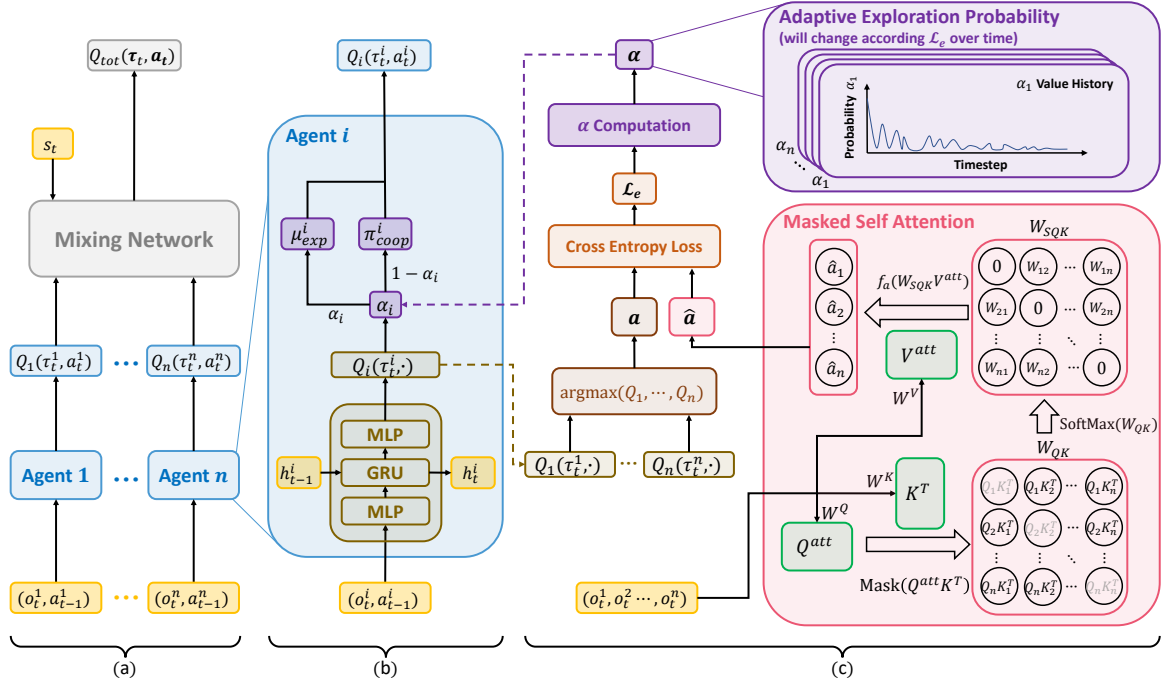


Figure 2: SMMAE framework. (a) Overall structure. (b) Agent network structure using α_i as the exploration probability and μ_{exp}^i as the exploration policy. (c) Network structure to compute adaptive exploration probability α .

where $q_{\xi_i}(a_i | o_{-i})$ is a variational posterior estimator of $p(a_i | o_{-i})$ with parameter ξ_i . As $p(a_i)$ is the prior of actions, the last term is a constant and can be ignored. Therefore, we only need to consider the first term. We use a network to represent $q_{\xi_i}(a_i | o_{-i})$, where the input is o_{-i} and the output is \hat{a}_i . Then the first term $\mathbb{E}_{a_i \sim p(a_i | o_{-i})} [\log q_{\xi_i}(a_i | o_{-i})]$ is the network's negative cross entropy loss $\mathcal{L}_{ce}(a_i, \hat{a}_i | o_{-i})$.

Based on this, we use a two-level heuristic method to adjust α_i :

$$\alpha_i^{new} = \begin{cases} \alpha_{low}, & \mathcal{L}_{ce}(a_i, \hat{a}_i | o_{-i}) \geq \mathcal{L}_{threshold} \\ \min\left(\alpha_i^{old} + \frac{\alpha_{high} - \alpha_{low}}{\lambda_\alpha}, \alpha_{high}\right), & \text{else,} \end{cases} \quad (4)$$

where α_{low} and α_{high} are the lower bound value and upper bound value, respectively. λ_α is the scaling parameter for α increasing steps and it takes λ_α steps to increase from the lower bound to the upper bound, and $\mathcal{L}_{threshold}$ is the loss threshold. For the policy convergence, α_{high} will decrease linearly to α_{low} at the end of the training. For stability, we update α every E_α episodes.

We need to mask agent i 's observation as we use other observation o_{-i} of agent i to predict agent i 's action. It means computing $\mathcal{L}_{ce}(a_i, \hat{a}_i | o_{-i})$ for each agent i requires optimizing n networks simultaneously and will consume large amounts of computing resources. Instead, we utilize the structure of the attention mechanism [41] to estimate all the n losses within one network:

$$g^{att}(o_t) = \text{SoftMax}\left(\frac{\text{Mask}\left(Q_t^{att}(K_t^{att})^T\right)}{\sqrt{d_k}}\right)V_t^{att}, \quad (5)$$

where $Q_t^{att} = o_t W^Q$, $K_t^{att} = o_t W^K$, $V_t^{att} = o_t W^V$ and W^Q, W^K, W^V denote the parameter matrix for the attention mechanism. Here d_k is the attention key size. As we use other observation o_{-i} of agent

i to predict, the function $\text{Mask}(X)$ will mask the correlation matrix and set the diagonal of the matrix to 0 (Figure 2). And the action prediction is: $\hat{a}_t = f_a(g^{att}(o_t))$, where f_a is an action prediction network with three fully-connected layers.

4.2 Explore by Maximizing State Entropy

Existing algorithms only focus on the global exploration policy, which will face the curse of dimensionality. Instead, we propose to use individual exploration and each agent can maximize its own exploration space, which is more efficient. SMMAE learns n independent exploration policies $\mu_{exp} \equiv \langle \mu_{exp}^1, \dots, \mu_{exp}^n \rangle$ for agents. Inspired by SMM [18], we use state marginal matching to maximize individual exploration space. In MARL, each agent only has access to local observation, so we use a local observation to approximate a local state and match the visited observation distribution $\rho_{\pi_i}(o^i)$ with a target distribution $p^*(o^i)$, where

$$\begin{aligned} & \rho_{\pi_i}(o^i) \\ &= \mathbb{E}_{a_t^i \sim \pi_i(a_t^i | o_t^i), o_{t+1}^i \sim O(P(s_{t+1} | s_t, a_t))} \left[\frac{1}{T} \sum_{t=1}^T \mathbf{1}(o_t^i = o^i) \right], \end{aligned} \quad (6)$$

and $p^*(o^i)$ is obtained using prior information. As there is not any prior information in our experiments, the target distribution of the exploration policy is uniform. To match the two distributions, for each policy we aim to optimize the following objective:

$$\begin{aligned} & \min_{\mu_{exp}^i} D_{KL}\left(\rho_{\mu_{exp}^i}(o^i) \| p^*(o^i)\right) \\ &= \max_{\mu_{exp}^i} \mathbb{E}_{o^i \sim \rho_{\mu_{exp}^i}} [\log p^*(o^i)] + \mathcal{H}_{\mu_{exp}^i}(o^i). \end{aligned} \quad (7)$$

As $p^*(o^i)$ is uniform, to optimize the objective, each exploration policy only needs to maximize the last term $\mathcal{H}_{\mu_{exp}^i}(o^i)$ and uses

Algorithm 1: SMMAE

Init: exploration probability α , exploration policy μ_{exp} , multi-agent Q-learning cooperation policy π_{coop} , exploration replay buffer \mathcal{D}_{exp} , cooperation replay buffer \mathcal{D}_{coop}

```

1 for episode = 1,  $\dots$ ,  $E$  do
2   Update  $\alpha$  every  $E_\alpha$  episodes according to Eq. 4
3   Reset the environment. Get state  $s_1$  and observations
    $\mathbf{o}_1 = (o_1^1, \dots, o_1^n)$ 
4   for  $t = 1, \dots, T$  do
5     for  $i = 1, \dots, n$  do
6       if  $x \sim \mathcal{U}(0, 1) < \alpha_i$  then
7         Use exploration policy to select action
          $a_t^i \sim \mu_{exp}^i(\tau_t^i)$ 
8       else
9         Use cooperation policy to select action
          $a_t^i \sim \pi_{coop}^i(\tau_t^i)$ 
10      end
11     end
12     Set joint action  $\mathbf{a}_t = (a_t^1, \dots, a_t^n)$ 
13     The environment takes a step. Get  $r_{env}, s_{t+1}, \mathbf{o}_{t+1}$ 
14     Compute  $\mathbf{r}_{exp}$  according to Eq. 8
15     Add transition tuples
        $\{(s_t, \mathbf{o}_t^i, a_t^i, s_{t+1}, \mathbf{o}_{t+1}^i, r_{exp}^i) \mid i = 1, \dots, n\}$  to  $\mathcal{D}_{exp}$ 
16     Add transition tuple  $\{(s_t, \mathbf{o}_t, \mathbf{a}_t, s_{t+1}, \mathbf{o}_{t+1}, r_{env})\}$  to
        $\mathcal{D}_{coop}$ 
17   end
18   Update density models  $q_{\xi_{exp}}^1, \dots, q_{\xi_{exp}}^n$  using
      $(\mathbf{o}_1, \dots, \mathbf{o}_T)$ 
19   Train exploration policy  $\mu_{exp}$  using  $\mathcal{D}_{exp}$ 
20   Train cooperation policy  $\pi_{coop}$  using  $\mathcal{D}_{coop}$ 
21 end

```

the following reward

$$r_{exp}(\mathbf{o}_t, \mathbf{a}_t) \triangleq \lambda_{exp} r_{env} - \log q_{\xi_{exp}}(\mathbf{o}_{t+1}) \quad (8)$$

to learn the exploration policy, where $\mathbf{r}_{env} \equiv \langle r_{env}, \dots, r_{env} \rangle$ and λ_{exp} is the scaling parameter for environment reward. We use environment reward here to assist the exploration. Here $\mathbf{r}_{exp}(\mathbf{o}_t, \mathbf{a}_t) \equiv \langle r_{exp}^1(o_t^1, a_t^1), \dots, r_{exp}^n(o_t^n, a_t^n) \rangle$, and the reward for the exploration policy of agent i is

$$r_{exp}^i(o_t^i, a_t^i) \triangleq \lambda_{exp} r_{env} - \log q_{\xi_{exp}}^i(o_{t+1}^i), \quad (9)$$

where $q_{\xi_{exp}}^i(o_{t+1}^i)$ is the variational estimator for the visiting probability of agent i 's observation o^i w.r.t $\rho_{\mu_{exp}^i}(o^i)$. In practice, each $q_{\xi_{exp}}^i(o_{t+1}^i)$ is estimated by a Variational Auto-Encoder (VAE) [15] model, where the input and the reconstruction objective are both o^i . The reconstruction loss of the VAE model is used as the last term $-\log q_{\xi_{exp}}^i(o_{t+1}^i)$.

We use the buffer \mathcal{D}_{coop} for multi-agent Q-learning and an extra buffer \mathcal{D}_{exp} for exploration policy learning to record the tuple $\{(s_t, o_t^i, a_t^i, s_{t+1}, o_{t+1}^i, r_{exp}^i) \mid i = 1, \dots, n\}$, where the trajectories in \mathcal{D}_{exp} are the same as that in \mathcal{D}_{coop} . The TD loss for all exploration policies μ_{exp} is:

$$\begin{aligned} \mathcal{L}_{exp} &= \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{exp}^i(\theta_{exp}^i) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(\tau_t^i, a_t^i, r_t^i, \tau_{t+1}^i) \sim \mathcal{D}_{exp}} \left[\left(r_t^i \right. \right. \\ &\quad \left. \left. + \gamma_{exp} \max_a Q_{exp}(\tau_{t+1}^i, a; \theta_{exp}^i) - Q_{exp}(\tau_t^i, a_t^i; \theta_{exp}^i) \right)^2 \right], \end{aligned} \quad (10)$$

where μ_{exp}^i is the exploration policy of agent i , θ_{exp}^i are the parameters of the target network and will be updated by θ_{exp}^i periodically. All the exploration policy μ_{exp} will be trained after the multi-agent Q-learning policy π_{coop} is trained.

The whole algorithm is summarized by Algorithm 1. We induce the adaptive exploration probability based on the uncertainty of the multi-agent system, and then adopt a new individual exploration policy. We first describe the adaptive way to control the exploration probability, which uses the correlation between action and observation of the agents as the criterion (lines 2 and 4-11). Then we introduce the individual exploration policy, which is optimized by state marginal matching [18] (lines 14-15 and 18-19).

5 EXPERIMENTS

We conduct experiments in this section to validate the efficiency of SMMAE, and we benchmark it on QMIX [34] for its widely proved coordination ability in MARL¹. Implementation details of SMMAE are shown in Appendix B. We first select a training process of SMMAE and study the effects of adaptive exploration probability on the observation areas of the agents during this training process. Then, we choose the benchmark SMAC [35] as the testbed to study the efficient exploration of SMMAE in complex MARL tasks.² We compare SMMAE with several methods, including EMC [53] that focuses on multi-agent exploration and some other baselines [21, 34, 36, 43, 44]. After that, we conduct ablation experiments to elaborate the effectiveness of each module. We also show the efficient exploration ability of SMMAE in a visual way. Finally, we show SMMAE is a general framework for MARL algorithms, and it can be applied for other environments like Level Based Foraging (LBF) [7]. More experimental results can be found in the appendix part.

5.1 Study on Adaptive Exploration

In this section, we analyze the effectiveness of adaptive exploration probability. To study the impact of adaptive exploration probability on one agent and the whole team, we experiment on the super hard map *6h_vs_8z* in SMAC, and a training process is shown in Figure 3.

The *left* figure of Figure 3 shows the adaptive exploration probability of one agent and the team test win rate changing over time. We find that the exploration probability of this agent shows a fluctuating decrease (the *blue* curve) as the training proceeds, which

¹The experiments are based on PyMARL framework.

²The SMAC version in our experiments is SC2.4.6.2.69232.

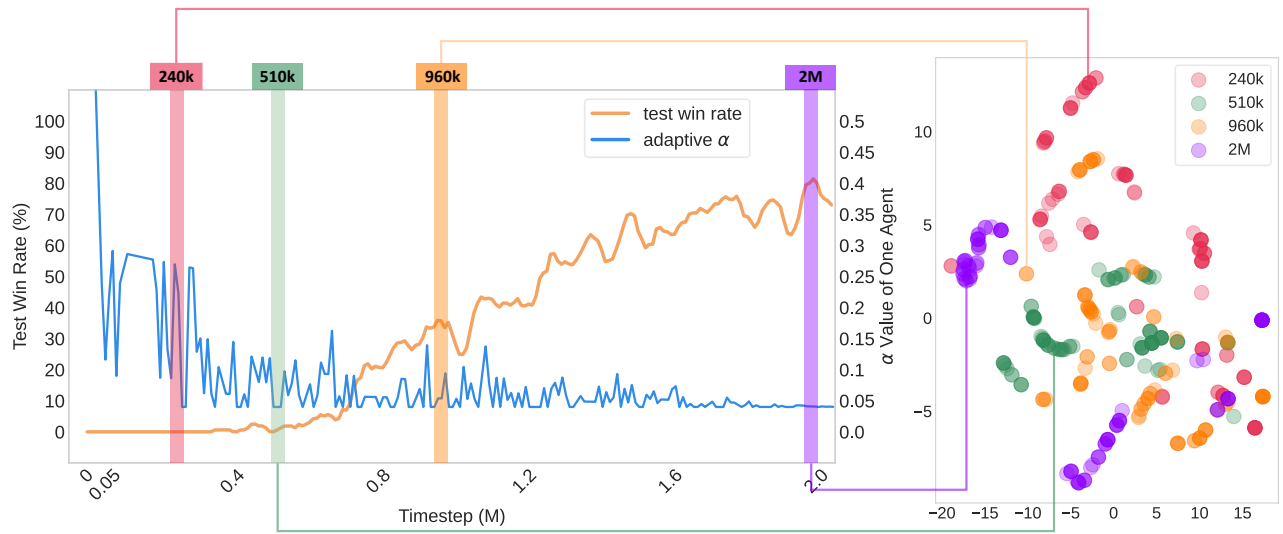


Figure 3: Influence of adaptive exploration probability on *6h_vs_8z*. The *left* part shows the correlation between the performance curve of self team and the adaptive α curve of one agent during 2 million training steps. The *right* part visualizes the agent’s observations (t-SNE projection on 2-D) during the training process, which illustrates the dynamic exploration preference and its impact on the performance.

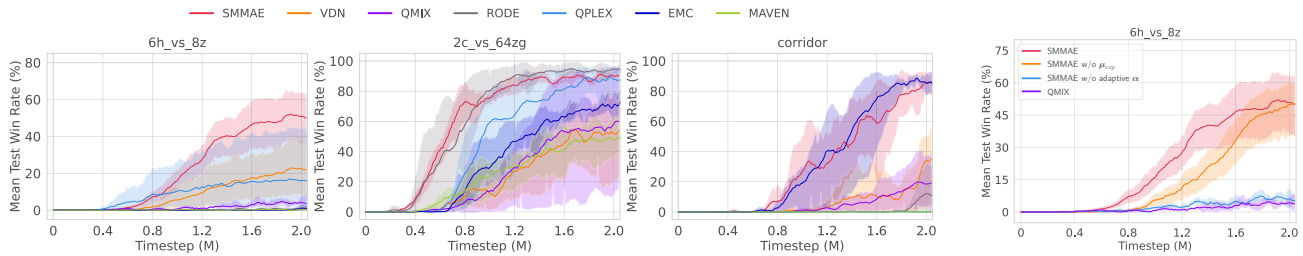


Figure 4: Results on three super hard maps in the SMAC.

Figure 5: Ablation study.

is adaptive compared with the vanilla ϵ -greedy. At the beginning of training, the agent explores with a high probability most of the time, while in the middle of training, the exploration probability of the agent varies continuously between high and low values. At the end of the training, the exploration probability tends to converge and gradually stabilizes at a low value.

To analyse the effect of exploration probability on observation areas, we select 4 timesteps, 240k, 510k, 960k, and 2M, where the exploration probability at 240k and 510k timesteps are the two peaks, while the exploration probability at 510k and 2M timesteps are the two valleys. All observations of this agent in the 10 episodes near these 4 timesteps are sampled (four shadow areas in the *left* figure of Figure 3), and 100 points were randomly selected from them. The *right* figure in Figure 3 shows the visualization of these points after we apply t-SNE [39] for dimensionality reduction. We find that for the two timesteps with higher exploration probability, there are more scattered clusters (240 and 960k, *red* and *orange* for each), while for the two timesteps with lower exploration probability, there are fewer clusters (510k and 2M, *green* and *violet*).

In addition, we can find some correlation between the *brown* test win rate curve and the *blue* exploration probability curve in the *left*

figure of Figure 3. The *red* shadow area in the figure is in the early stages of training, where the test win rate remains 0% and the agent focuses on exploring and experiencing more state areas. The *green* shadow area in the figure has a low probability of exploration. At these timesteps, the correlation between the agents is low, so the agents are more focused on learning how to cooperate with other agents than exploring. Then, some improvement in the *brown* test win curve can be noticed in this region. The *orange* shadow area in the figure is a peak of exploration. At this time, the correlation between the agents is already high, so the agent learns to explore more areas, thus jumping out of the local optimum. Therefore, there is a significant drop in the test win rate near this area. In the *violet* region, the probability of exploration is close to convergence and the agent focuses on better cooperation. The test win rate reaches the peak of the entire training process, validating the effectiveness of our design of adaptive exploring probability.

5.2 Performance on SMAC Super Hard Maps

In order to test the effectiveness of SMMAE, we use the benchmark SMAC [35] as the testbed. We compare SMMAE with some baselines, where VDN [36] and QMIX [34] are popular baselines, RODE [44] and QPLEX [43] are the state-of-the-art baselines, and

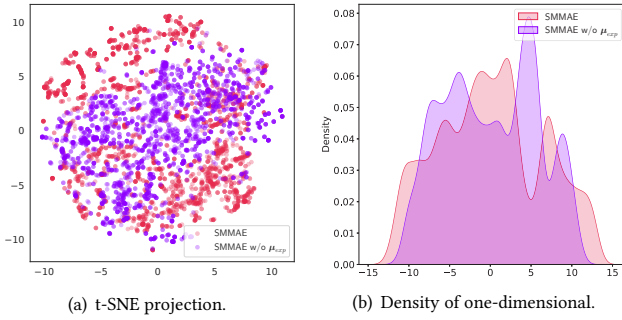


Figure 6: The t-SNE projection on global state.

MAVEN [21] and EMC [53] are the latest methods on MARL exploration. For evaluation, we carry out each experiment with 5 random seeds, and the results are shown with a 95% confidence interval.

Figure 4 shows the training curves of SMMAE and other methods on several super hard tasks of SMAC. It can be seen that SMMAE achieves competitive results on tasks such as *6h_vs_8z* and *corridor* that require sufficient individual exploration and team cooperation. It shows that adaptive exploration can also achieve the results of other methods using complex modules. This also illustrates the effectiveness of SMMAE because it performs individual exploration and task-related exploration (Section 5.5 gives a visual explanation).

The final test win rate of SMMAE and that of baseline EMC on *corridor* are almost the same, but it should be noted that EMC needs to use 15GB of GPU when training *corridor*, while our SMMAE only needs about 4GB.

5.3 Ablation Study

As our method includes multiple modules, we compare SMMAE with three ablated methods for ablation study:

- SMMAE w/o μ_{exp} : SMMAE using random exploration instead of the exploration policy μ_{exp} .
- SMMAE w/o adaptive α : SMMAE using fixed ending α instead of the adaptive α .
- QMIX: the QMIX baseline [34].

Figure 5 shows the average test win rate of these four approaches on the *6h_vs_8z* scenario of SMAC. Compared with the original QMIX, SMMAE w/o adaptive α is slightly better, while SMMAE w/o μ_{exp} has greatly improves performance. It illustrates the effectiveness of the two modules. Among these four methods, SMMAE achieves the best results, which indicates the effectiveness of using two modules simultaneously. However, it is also observed that the variance of SMMAE becomes larger compared to SMMAE w/o μ_{exp} , which indicates that increasing the exploration capability of the individual exploration policy may introduce some instability to the training process. Compared with SMMAE w/o μ_{exp} , SMMAE has similar convergence win rate. Therefore, We conducted experiments on two additional maps *2c_vs_64zg* and *corridor* in SMAC to further investigate the effect of the exploration policies μ_{exp} in SMMAE (Figure 7). All of these results demonstrate that although random exploration can sometimes perform well, μ_{exp} can explore more efficiently and speed up training.

To qualitatively analyze the exploration ability of the exploration policy, in the ablation experiment, we sample the episodes experienced by SMMAE and SMMAE w/o μ_{exp} . Then we sample global

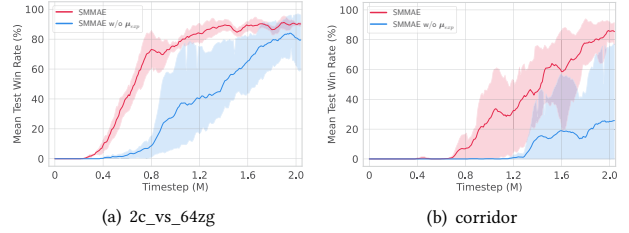


Figure 7: More exploration policy ablation in SMAC.

state points randomly from them, and visualize the points after dimensionality reduction using t-SNE [39]. Figure 6 shows the results. It can be found that SMMAE has more dispersed clusters in the two-dimensional t-SNE embedding space (Figure 6(a)), and SMMAE has more uniform density and wider range in one-dimensional t-SNE embedding space (Figure 6(b)). It shows that SMMAE, which uses the exploration policy μ_{exp} that maximizes the state entropy of local visited observations, has a stronger exploration ability. Then we try to express the visualization results quantitatively. Because the state space is continuous, we retain two decimal places for each dimension of each global state and calculate an approximate entropy of the visited state based on state counting. The entropy of SMMAE is 15.238, and the entropy of SMMAE w/o μ_{exp} is 14.982. The quantitative results are consistent with the visualization results (Figure 6).

5.4 Comparison with Finetuned-QMIX

Hu et al. [14] finetune QMIX [34] by adding a variety of implementation tricks and achieve good performance on SMAC. To show the necessity of the designs in SMMAE, we compare SMMAE with it, denoted as *Hu-QMIX*, on the four super hard maps in SMAC. The results show that our SMMAE, which is based on the vanilla QMIX outperforms the finetuned QMIX *Hu-QMIX* on all the maps (Figure 9). Compared with *Hu-QMIX*, SMMAE is able to achieve higher test win rate faster and converge to higher win rate on all maps. It demonstrates the effectiveness of SMMAE.

5.5 Efficient Exploration in Task-Related Space

In this section, we qualitatively compare the exploration efficiency of SMMAE with that of baseline EMC on *6h_vs_8z* by visualization. Figure 8 shows the visual results. The *left* picture and the *right* figure are the coordination patterns of the two algorithms during test time. Same as the operation in Figure 6(a), we sample the global state points from their training processes and use t-SNE [39] to visualize the points after dimensionality reduction (*middle* figure of Figure 8). The *middle* figure of Figure 8 shows that SMMAE has a similar number of clusters as that of EMC, which means SMMAE and EMC have similar sizes of exploration areas. Same as the calculation method of visiting state entropy in Section 5.3, the entropy of SMMAE is 13.574, and the entropy of EMC is 13.428. The replays of SMMAE during test time show that the agents have learned how to cooperate in different ways to attack the enemies, where the agents have learned how to surround the enemies and how to increase the distance to the enemies to take advantage of range (*left* figure of Figure 8). However, the replays of EMC during test time show that the agents only have learned how to escape (*right* figure of Figure 8). The different results demonstrate

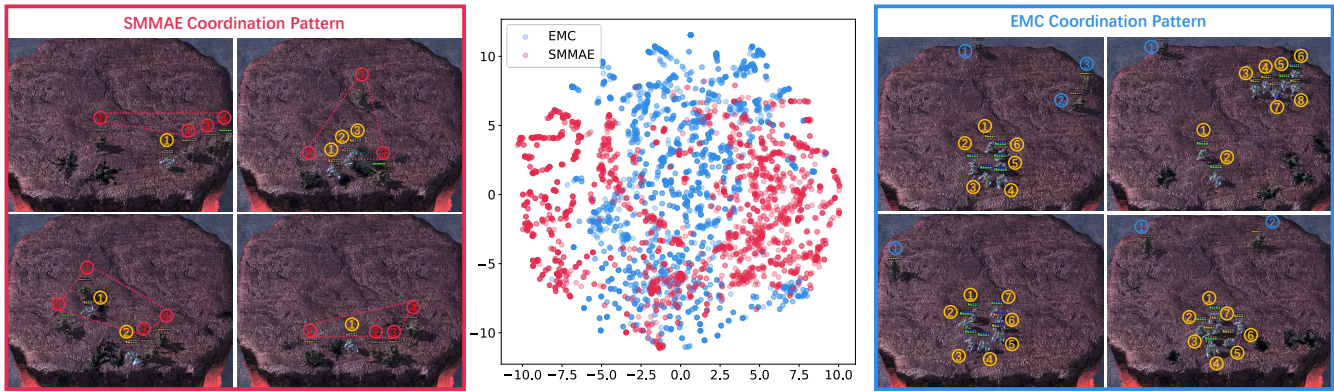


Figure 8: Exploration experiments on *6h_vs_8z*. *Left*: SMMAE replay during test, where the red ones are controlled by SMMAE and the orange ones are the enemies. *Middle*: Two-dimensional global state t-SNE embeddings during training. *Right*: EMC replay during test, where the blue ones are controlled by EMC and the orange ones are the enemies.

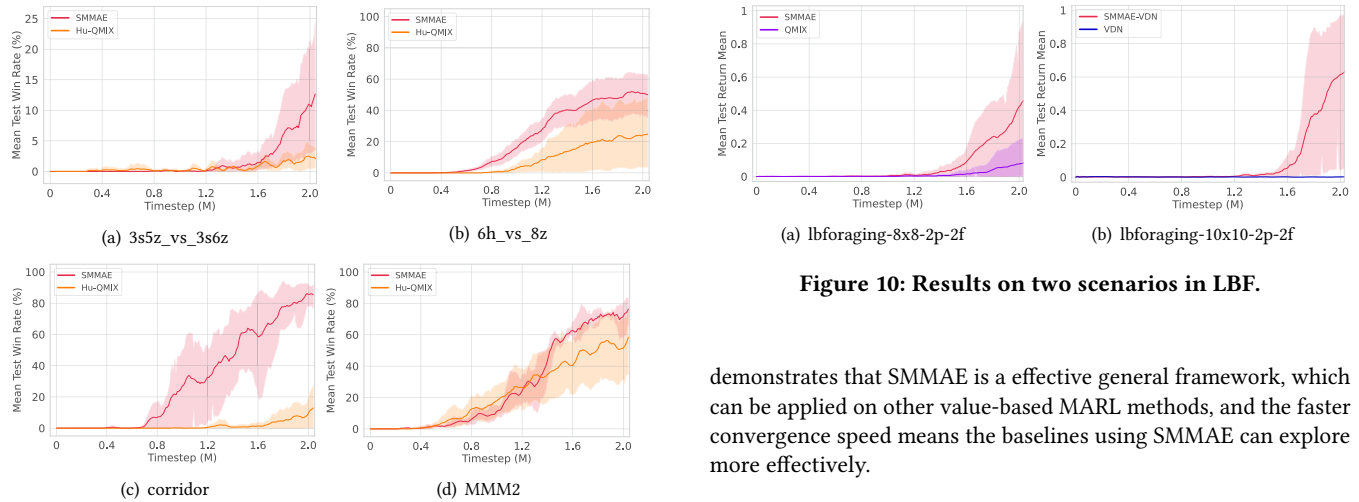


Figure 9: Comparison with the finetuned QMIX on four super hard maps of SMAC.

that although SMMAE and EMC have similar sizes of exploration areas, SMMAE can explore more efficiently because the space it explores is more task-related. In fact, our method itself does not explicitly aim to explore the task-related space. In SMMAE, each agent uses its own exploration policy to maximize the entropy of the states it visits, so the method can more efficiently cover a wider exploration space, thus covering task-related states.

5.6 Performance on Level Based Foraging and Application on VDN

Besides SMAC, we also use Level Based Foraging (LBF) [7] to test the ability of SMMAE (Figure 10). The results show that SMMAE can outperform QMIX on the simple task (Figure 10(a)). Because VDN performs better than QMIX on LBF using PyMARL, we also apply SMMAE to another baseline VDN, denoted as SMMAE-VDN, to test the ability of SMMAE. The results show that SMMAE-VDN can also outperform VDN on the task in LBF (Figure 10(b)). It

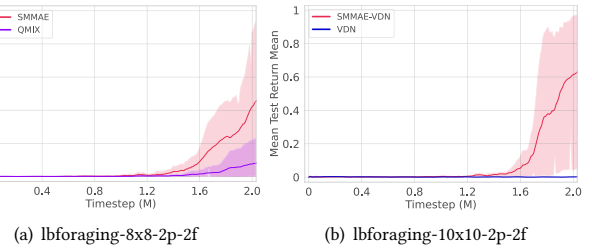


Figure 10: Results on two scenarios in LBF.

demonstrates that SMMAE is an effective general framework, which can be applied on other value-based MARL methods, and the faster convergence speed means the baselines using SMMAE can explore more effectively.

6 CONCLUSION

With appropriate hyper-parameters, we empirically find that vanilla MARL algorithms using self-exploration can also achieve competitive performance. Consequently, in this paper we propose SMMAE, a novel algorithm that adaptively adjusts the individual exploration probability according to the uncertainty of the multi-agent system at different timesteps. SMMAE focuses on each agent’s individual exploration ability by learning an individual exploration policy, which is optimized by target state distribution matching. We study SMMAE on a variety of tasks in the SMAC benchmark, and empirically reveal that SMMAE can explore more efficiently on task-related states and generate better cooperation policies. We take a step towards achieving a trade-off between individual exploration and team cooperation, and we think it is promising for SMMAE to solve complex tasks in multi-agent systems combined with other methods like communication in future work. It should be pointed out that directly matching the explored states to a uniform distribution may cause unexpected exploration, and the adaptive exploration probability is necessary for stabilizing training. It deserves further research on the target distribution of the exploration policy and the adaptive exploration probability.

7 ACKNOWLEDGMENTS

This work is supported by National Key Research and Development Program of China (2020AAA0107200), the National Science Foundation of China (62250069, 61921006), and the program B for Outstanding Ph.D. candidate of Nanjing University. We would like to thank Shenghua Wan, and the anonymous reviewers for their helpful discussions and support.

REFERENCES

- [1] Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturovski, Pablo Sprechmann, Alex Vitvitskiy, Zhaohan Daniel Guo, and Charles Blundell. 2020. Agent57: Outperforming the atari human benchmark. In *ICML*. 507–517.
- [2] Adrià Puigdomènech Badia, Pablo Sprechmann, Alex Vitvitskiy, Daniel Guo, Bilal Piot, Steven Kapturovski, Olivier Tieleman, Martin Arjovsky, Alexander Pritzel, et al. 2020. Never give up: Learning directed exploration strategies. In *ICLR*.
- [3] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. Unifying count-based exploration and intrinsic motivation. In *NeurIPS*.
- [4] Glen Berseth, Daniel Geng, Coline Devin, Nicholas Rhinehart, Chelsea Finn, Dinesh Jayaraman, and Sergey Levine. 2021. SMiRL: Surprise Minimizing Reinforcement Learning in Unstable Environments. In *ICLR*.
- [5] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. 2019. Exploration by random network distillation. In *ICLR*.
- [6] Jiahao Cao, Lei Yuan, Jianhao Wang, Shaowei Zhang, Chongjie Zhang, Yang Yu, and De-Chuan Zhan. 2021. LINDA: Multi-Agent Local Information Decomposition for Awareness of Teammates. *arXiv preprint arXiv:2109.12508* (2021).
- [7] Filippos Christianos, Lukas Schäfer, and Stefano Albrecht. 2020. Shared experience actor-critic for multi-agent reinforcement learning. *NeurIPS* (2020).
- [8] Will Dabney, Georg Ostrovski, and André Barreto. 2021. Temporally-Extended ϵ -Greedy Exploration. In *ICLR*.
- [9] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. 2019. Diversity is all you need: Learning skills without a reward function. In *ICLR*.
- [10] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual multi-agent policy gradients. In *AAAI*.
- [11] Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Matteo Hessel, Ian Osband, Alex Graves, Volodymyr Mnih, Rémi Munos, Demis Hassabis, Olivier Pietquin, Charles Blundell, and Shane Legg. 2018. Noisy Networks For Exploration. In *ICLR*.
- [12] Nir Greshler, Ofir Gordon, Oren Salzman, and Nahum Shimkin. 2021. Cooperative Multi-Agent Path Finding: Beyond Path Planning and Collision Avoidance. In *MRS*. 20–28.
- [13] Sven Gronauer and Klaus Diepold. 2022. Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review* 55, 2 (2022), 895–943.
- [14] Jian Hu, Siyang Jiang, Seth Austin Harding, Haibin Wu, and Shih-wei Liao. 2021. Rethinking the implementation tricks and monotonicity constraint in cooperative multi-agent reinforcement learning. *CoRR* (2021).
- [15] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- [16] Gary A Klein, Judith Orasanu, Roberta Calderwood, Caroline E Zsombok, et al. 1993. *Decision Making in Action: Models and Methods*. Ablex Norwood, NJ.
- [17] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. 2004. Estimating mutual information. *Physical review E* 69, 6 (2004), 066138.
- [18] Lisa Lee, Benjamin Eysenbach, Emilio Parisotto, Eric Xing, Sergey Levine, and Ruslan Salakhutdinov. 2019. Efficient exploration via state marginal matching. *CoRR abs/1906.05274* (2019).
- [19] Iou-Jen Liu, Unnat Jain, Raymond A Yeh, and Alexander Schwing. 2021. Cooperative exploration for multi-agent deep reinforcement learning. In *ICML*. 6826–6836.
- [20] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *NeurIPS*.
- [21] Anuj Mahajan, Tabish Rashid, Mikayel Samvelyan, and Shimon Whiteson. 2019. Maven: Multi-agent variational exploration. *NeurIPS* (2019).
- [22] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [23] Frans A Oliehoek and Christopher Amato. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer.
- [24] Frans A Oliehoek, Matthijs TJ Spaan, and Nikos Vlassis. 2008. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research* 32 (2008), 289–353.
- [25] Afshin Oroojlooyjadid and Davood Hajinezhad. 2019. A review of cooperative multi-agent deep reinforcement learning. *arXiv preprint arXiv:1908.03963* (2019).
- [26] Georg Ostrovski, Marc G Bellemare, Aaron Oord, and Rémi Munos. 2017. Count-based exploration with neural density models. In *ICML*. 2721–2730.
- [27] Simone Parisi, Victoria Dean, Deepak Pathak, et al. 2021. Interesting Object, Curious Agent: Learning Task-Agnostic Exploration. *NeurIPS* (2021).
- [28] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *ICML*. 2778–2787.
- [29] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. 2019. Self-supervised exploration via disagreement. In *ICML*. 5062–5071.
- [30] Miruna Pislari, David Szepesvari, Georg Ostrovski, Diana Borsa, and Tom Schaul. 2022. When should agents explore?. In *ICLR*.
- [31] Silviu Pitis, Harris Chan, Stephen Zhao, Bradly Stadie, and Jimmy Ba. 2020. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *ICML*. 7750–7761.
- [32] Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y. Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. 2018. Parameter Space Noise for Exploration. In *ICLR*.
- [33] Tabish Rashid, Gregory Farquhar, Bei Peng, and Shimon Whiteson. 2020. Weighted qmix: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning. *NeurIPS* (2020), 10199–10210.
- [34] 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 *ICML*. 4295–4304.
- [35] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder De Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. 2019. The starcraft multi-agent challenge. In *AAMAS*. 2186–2188.
- [36] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 2018. Value-decomposition networks for cooperative multi-agent learning. In *AAMAS*. 2085–2087.
- [37] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [38] Michel Tokic. 2010. Adaptive ϵ -greedy exploration in reinforcement learning based on value differences. In *AAAI*. 203–210.
- [39] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, 11 (2008), 2579–2605.
- [40] Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. In *AAAI*.
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.
- [42] Huimin Wang and Kam-Fai Wong. 2021. A Collaborative Multi-agent Reinforcement Learning Framework for Dialog Action Decomposition. In *EMNLP*. 7882–7889.
- [43] Jianhao Wang, Zhizhou Ren, Terry Liu, Yang Yu, and Chongjie Zhang. 2021. QPLEX: Duplex Dueling Multi-Agent Q-Learning. In *ICLR*.
- [44] Tonghan Wang, Tarun Gupta, Anuj Mahajan, Bei Peng, Shimon Whiteson, and Chongjie Zhang. 2021. RODE: Learning Roles to Decompose Multi-Agent Tasks. In *ICLR*.
- [45] Tonghan Wang, Jianhao Wang, Yi Wu, and Chongjie Zhang. 2020. Influence-based multi-agent exploration. In *ICLR*.
- [46] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. 2016. Dueling network architectures for deep reinforcement learning. In *ICML*. 1995–2003.
- [47] Ke Xue, Jiacheng Xu, Lei Yuan, Miqing Li, Chao Qian, Zongzhang Zhang, and Yang Yu. 2022. Multi-agent Dynamic Algorithm Configuration. In *NeurIPS*.
- [48] Tianpei Yang, Hongyao Tang, Chenjia Bai, Jinyi Liu, Jianye Hao, Zhaopeng Meng, and Peng Liu. 2021. Exploration in Deep Reinforcement Learning: A Comprehensive Survey. *CoRR abs/2109.06668* (2021).
- [49] Chao Yu, Akash Velu, Eugene Vinytsky, Yu Wang, Alexandre Bayen, and Yi Wu. 2021. The surprising effectiveness of ppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955* (2021).
- [50] Lei Yuan, Chenghe Wang, Jianhao Wang, Fuxiang Zhang, Feng Chen, Cong Guan, Zongzhang Zhang, Chongjie Zhang, and Yang Yu. 2022. Multi-Agent Concentrative Coordination with Decentralized Task Representation. In *IJCAI*. 599–605.
- [51] Lei Yuan, Jianhao Wang, Fuxiang Zhang, Chenghe Wang, Zongzhang Zhang, Yang Yu, and Chongjie Zhang. 2022. Multi-Agent Incentive Communication via Decentralized Teammate Modeling. In *AAAI*. 9466–9474.
- [52] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E Gonzalez, and Yuandong Tian. 2021. NovelD: A Simple yet Effective Exploration Criterion. In *NeurIPS*.
- [53] Lulu Zheng, Jiarui Chen, Jianhao Wang, Jiamin He, Yujing Hu, Yingfeng Chen, Changjie Fan, Yang Gao, and Chongjie Zhang. 2021. Episodic Multi-agent Reinforcement Learning with Curiosity-driven Exploration. In *NeurIPS*.
- [54] Ming Zhou, Jun Luo, Julian Vilella, Yaodong Yang, David Rusu, Jiayu Miao, Weinan Zhang, Montgomery Alban, Iman Fadarar, Zheng Chen, et al. 2020. Smarts: Scalable multi-agent reinforcement learning training school for autonomous driving. *CoRR abs/2010.09776* (2020).