

MESA: Cooperative Meta-Exploration in Multi-Agent Learning through Exploiting State-Action Space Structure

Zhicheng Zhang*

Carnegie Mellon University
Pittsburgh, Pennsylvania, United States
zhichen3@cs.cmu.edu

Yi Wu

Tsinghua University
Beijing, China
jxwuyi@gmail.com

Yancheng Liang*

University of Washington
Seattle, Washington, United States
yancheng@cs.washington.edu

Fei Fang

Carnegie Mellon University
Pittsburgh, Pennsylvania, United States
feif@cs.cmu.edu

ABSTRACT

Multi-agent reinforcement learning (MARL) algorithms often struggle to find strategies close to Pareto optimal Nash Equilibrium, owing largely to the lack of efficient exploration. The problem is exacerbated in sparse-reward settings, caused by the larger variance exhibited in policy learning. This paper introduces MESA, a novel meta-exploration method for cooperative multi-agent learning. It learns to explore by first identifying the agents' high-rewarding joint state-action subspace from training tasks and then learning a set of diverse exploration policies to "cover" the subspace. These trained exploration policies can be integrated with any off-policy MARL algorithm for test-time tasks. We first showcase MESA's advantage in a multi-step matrix game. Furthermore, experiments show that with learned exploration policies, MESA achieves significantly better performance in sparse-reward tasks in several multi-agent particle environments and multi-agent MuJoCo environments, and exhibits the ability to generalize to more challenging tasks at test time.

KEYWORDS

Multi-Agent Reinforcement Learning; Meta-Learning; Exploration Strategy

ACM Reference Format:

Zhicheng Zhang, Yancheng Liang, Yi Wu, and Fei Fang. 2024. MESA: Cooperative Meta-Exploration in Multi-Agent Learning through Exploiting State-Action Space Structure. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

Reinforcement learning (RL) algorithms often adopt a trial-and-error learning paradigm and optimize the policy based on the reward signals given by the environment. The effectiveness of RL relies on efficient exploration, especially in sparse reward settings,

*Equal contribution.



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

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

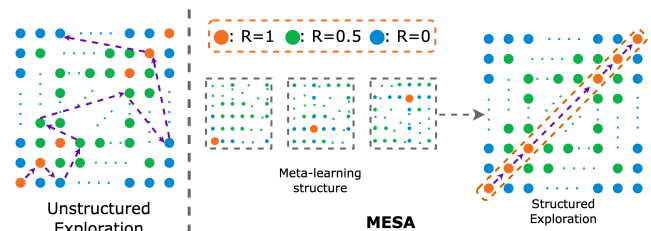


Figure 1: Illustration of structured exploration and unstructured exploration behavior in the 2-player climb game. The rows and columns indicate the players' action space. While unstructured exploration aims to visit novel states, structured exploration exploits structures in the joint state-action space, helping agents coordinately and more efficiently explore the potential high-reward subspace.

as it is critical to get sufficient experiences with high rewards to guide the training.

The exploration challenge has been studied extensively and existing works can be categorized mainly into two streams. One core idea with great success is to incentivize the agent to visit the under-explored states more frequently by adding an intrinsic reward based on a visitation measure [3, 25, 28, 37] or some other heuristics [17, 39].

However, in multi-agent settings, due to the exponential growth of the joint state-action space, simply visiting more novel states can be increasingly ineffective. Exploration policies need to better capture the low-dimensional structure of the tasks and leverage the structural knowledge for higher exploration efficiency.

Another line of work specifically learns exploration strategies. However, these works do not explicitly consider the underlying task structure. For example, Mahajan et al. conditions the policy on a shared latent variable [24] learned via mutual information maximization. Liu et al. adopts a goal-conditioned exploration strategy by setting state features as goals [21]. Other works in the single-agent settings [6, 26, 35] learn exploration policies through a pre-defined intrinsic reward. All these works train the exploration policy using task-agnostic exploration-specific rewards. In

Section 4, we will present a simple matrix game to show that popular exploration methods can have difficulties finding the optimal solution due to the reward structure of the game.

How can we enable the agents to more effectively explore by leveraging the intrinsic structure of the environment? We adopt a meta-exploration framework (i.e., learning to explore) for MARL: we first train multiple *structured exploration policies* from a set of training tasks (referred to as the meta-training stage), and then use these exploration policies to facilitate agents’ learning in a test-time task, which is typically a new task sampled from the task distribution (referred to as meta-testing stage). We develop a multi-agent meta-exploration method, *Cooperative Meta-Exploration in Multi-Agent Learning through Exploiting State-Action Space Structure* (MESA) for fully cooperative settings. MESA leverages the task structures by explicitly identifying the agents’ high-rewarding joint state-action subspace in the training tasks. It then trains a set of diverse exploration policies to cover this identified subspace. The exploration policies are trained with a reward scheme induced by the distance to the high-rewarding subspace. The meta-learned exploration policies can be combined with any off-policy MARL algorithm during the meta-testing stage by randomly selecting learned exploration policies to collect valuable experiences. Such structured exploration can help the agents to learn good joint policies efficiently (Figure 1). We empirically show the success of MESA on the matrix climb game and its harder multi-stage variant. In addition, we evaluate MESA in two continuous control tasks, i.e., the MPE environment [23] and the multi-agent MuJoCo benchmark [29]. We demonstrate the superior performance of MESA compared to existing multi-agent learning and exploration algorithms. Furthermore, we show that MESA is capable of generalizing to unseen test-time tasks that are more challenging than any of the training tasks.

2 RELATED WORK

Exploration has been a long-standing challenge in RL with remarkable progress achieved in the single-agent setting [3, 5, 10, 25, 28, 34, 37]. Most of these works maintain pseudo-counts over states and construct intrinsic rewards to encourage the agents to visit rarely visited states more frequently [3, 25, 28, 37]. These count-based methods have been extended to the multi-agent setting by incentivizing intra-agent interactions or social influence [17–19, 39]. However, in the multi-agent setting, a simple count-based method can be less effective due to the partial observability of each agent, an exponentially large joint state-action space, and the existence of multiple non-Pareto-optimal NE. Therefore, recent works focus on discovering the structures of possible multi-agent behaviors. For example, [24] adopts variational inference to learning structured latent-space-policies; [15] generates similar tasks with simpler reward functions to promote cooperation; [21] learns to select a subset of state dimensions for efficient exploration. We follow a meta-learning framework and learn structured exploration strategies by exploiting high-rewarding subspace in the joint state-action space. Our method also leverages a count-based technique as a subroutine during the meta-training phase to prevent over-exploitation and mode collapse.

Meta reinforcement learning (meta-RL) is a popular RL paradigm that focuses on training a policy that can quickly adapt on an unseen task at test time [9, 12, 14, 20, 32, 40, 42, 44]. Such a paradigm has been extended to the setting of learning to explore. The key idea is to meta-learn a separate exploration policy that can be used in the testing task. Most closely related to our work is [26], where an exploration policy is pretrained on a set of training tasks. However, their method is designed for the single-agent setting and learns the exploration policy by using a task-agnostic intrinsic reward to incentivize visitation of interesting states, while we directly utilize the task reward to learn the structure of the environments. Other existing works in meta-exploration propose to learn a latent-space exploration policy that is conditioned on a task variable, which can be accomplished by meta-policy gradient [14, 20, 40], variational inference [32] or information maximization [42] over the training tasks. Therefore, at test time, posterior inference can be performed for the latent variable towards fast exploration strategy adaption. Our approach follows a similar meta-exploration paradigm by learning additional exploration policies. However, existing meta-exploration methods focus on the single-agent setting while we consider much more challenging multi-agent games with a distribution of similarly-structured tasks, for example, the MPE environment [23] with a distribution of target landmarks that the agents need to reach. In addition, we meta-learn a discrete set of exploration policies through an iterative process, which results in a much simpler meta-testing phase without the need for posterior sampling or gradient updates on exploration policies. Besides, some other methods pretrain exploration policies from an offline dataset [7, 31, 36], which is beyond the scope of this paper.

Finally, our approach largely differs from the setting of multi-task learning [1, 2, 11, 16, 27], which are commonly evaluated in environments with heterogeneous tasks or scenarios. Our exploration policies are *not* trained to achieve high returns in the training tasks. Instead, they are trained to reach as many high-reward *state-action pairs* as possible collected in a diverse set of tasks. Therefore, the state-action pairs covered by a single exploration policy are very likely to be distributed across different training tasks.

3 PRELIMINARIES

Dec-POMDP. We consider fully-cooperative Markov games described by a decentralized partially observable Markov decision process (Dec-POMDP), which is defined by $\langle \mathcal{S}, \mathcal{A}, P, R, \Omega, O, n, \gamma \rangle$. \mathcal{S} is the state space. $\mathcal{A} \equiv \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ is the joint action space. The dynamics is defined by the transition function $P(s' | s, \mathbf{a})$. Agents share a reward function $R(s, \mathbf{a})$, and $\gamma \in (0, 1)$ is the discount factor. $\Omega \equiv \Omega_1 \times \dots \times \Omega_n$ is the joint observation space, where Ω_i is the observation space for agent i . At each timestep, each agent i only has access to its own observation $o_i \in \Omega_i$ defined by the function $O : \mathcal{S} \times \mathcal{A} \mapsto \Omega$. The goal of agents in Dec-POMDP is to maximize the common expected discounted return under the joint policy π : $\mathcal{J}(\pi) = \mathbb{E}_\pi \left[\sum_t \gamma^t R(s_t, \mathbf{a}_t) \right]$.

Learning to Explore. Meta-RL assumes a task distribution $p(\mathcal{T})$ over tasks, and an agent aims to learn to quickly adapt to a test-time task $\mathcal{T}_{\text{test}}$ drawn from $p(\mathcal{T})$ after training in a batch of training tasks $\{\mathcal{T}_i | \mathcal{T}_i \sim p(\mathcal{T})\}_{i=1}^B$. Inspired by the explicit exploration methods [6, 42], we adopt a meta-exploration framework for MARL: we

learn joint exploration policies π_e from training tasks $\{\mathcal{T}_i \mid \mathcal{T}_i \sim p(\mathcal{T})\}_{i=1}^B$ and use π_e to collect experiences for the training of the agents' policy profile π in task $\mathcal{T}_{\text{test}}$, denoted as $\pi(\pi_e, \mathcal{T}_{\text{test}})$. Formally, the objective of meta-exploration is

$$\max_{\pi_e} \mathbb{E}_{\mathcal{T}_{\text{test}} \sim p(\mathcal{T})} \left[\mathbb{E}_{\pi(\pi_e, \mathcal{T}_{\text{test}})} \left[\sum_t \gamma^t R_i(s_t, \mathbf{a}_t) \right] \right]. \quad (1)$$

Nash Equilibrium and Pareto Optimality. A joint policy π is an NE if each agent's policy π_i is a best response to the other agents' policies π_{-i} . That is, for any agent i 's alternative policy π'_i , we have $Q_i(\pi) \geq Q_i(\pi'_i, \pi_{-i})$, where Q_i is the value function for agent i . A joint policy π is Pareto optimal if there does not exist an alternative joint policy π' such that $\forall i, Q_i(\pi') \geq Q_i(\pi)$ and $\exists i, Q_i(\pi') > Q_i(\pi)$.

4 A MOTIVATING EXAMPLE: CLIMB GAME

We analyze a fully cooperative matrix game known as Climb Game. In Section 4.1, we show how popular exploration strategies, including unstructured strategies like uniform exploration and task-specific strategies like ϵ -greedy, fail to efficiently explore the climb game. By contrast, we show in Section 4.2 that a simple structured exploration strategy can substantially improve the exploration efficiency.

A climb game $G_f(n, u, U)$ is a n -player game with action space $\mathcal{A}_i = \{0, \dots, U-1\}$ for any player i . The reward of a joint action $\mathbf{a} \in \mathcal{A}$ is determined by the number of players performing a specific action u (denoted as $\#u$), which is

$$R(\mathbf{a}) = \begin{cases} 1, & \text{if } \#u = n, \\ 1 - \delta \ (0 < \delta < 1), & \text{if } \#u = 0, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

4.1 Exploration Challenge

A climb game $G_f(n, u, U)$ has three groups of NE: the Pareto optimal NE (u, u, \dots, u) , the sub-optimal NEs $\{(a_1, a_2, \dots, a_n) \mid \forall i, a_i \neq u\}$, and the zero-reward NEs $\{(a_1, a_2, \dots, a_n) \mid 1 < \#u < n\}$. The sheer difference in the size of the three subsets of NEs makes it particularly challenging for RL agents to learn the optimal policy profile without sufficient exploration, as evidenced by the theoretical analysis below and empirical evaluation in Section 6.

Consider a 2-agent climb game $G_f(2, 0, U)$. A joint action \mathbf{a} can be represented by a pair of one-hot vectors $[\mathbf{e}_i, \mathbf{e}_j] \in \{0, 1\}^{2U}$. Let $q(\mathbf{x}, \mathbf{y}; \theta)$ be a joint Q function parameterized by θ that takes input $\mathbf{x}, \mathbf{y} \in \{0, 1\}^U$ and is learned to approximate the reward of the game. We hope the joint Q function has the same optimal policy profile.

Definition 4.1. We call a joint Q function $q(\mathbf{x}, \mathbf{y}; \theta)$ *equivalently optimal* when $q(\mathbf{e}_0, \mathbf{e}_0; \theta) = \max_{0 \leq i, j < U} q(\mathbf{e}_i, \mathbf{e}_j; \theta)$. When a joint Q function is equivalently optimal, one can use it to find the optimal policy.

Since neural networks are difficult to analyze in general [4], we parameterize the joint Q function in a quadratic form:

$$q(\mathbf{x}, \mathbf{y}; \mathbf{W}, \mathbf{b}, \mathbf{c}, d) = \mathbf{x}^\top \mathbf{W} \mathbf{y} + \mathbf{b}^\top \mathbf{x} + \mathbf{c}^\top \mathbf{y} + d \quad (3)$$

A Gaussian prior $p(\mathbf{W}) = \mathcal{N}(\mathbf{W}; 0, \sigma_w^2 I)$ is introduced under the assumption that a non-linear \mathbf{W} is harder and slower to learn.

Quadratic functions have been used in RL [13, 38] as a replacement for the commonly-used multi-layer perceptron, and there are also theoretical results [8] analyzing neural networks with quadratic activation. For the climb game, it is easy to verify that the quadratic coefficients make the joint Q function sufficiently expressive to perfectly fit the reward function by setting \mathbf{W} to be the reward matrix. Therefore, the learning process of Q is mainly affected by how the exploration policy samples the data.

Consider an exploration policy $p_e^{(t)}$ that selects joint action $\mathbf{a} = (i, j)$ at step t with probability $p_e^{(t)}(i, j)$. The efficiency of an exploration policy can be measured by the required number of steps for learning an equivalently optimal Q function using the maximum likelihood estimator over the data sampled from $p_e^{(t)}$. The learning objective includes both the prior $p(\mathbf{W})$ and the likelihood of prediction error $p(E_{ij})$, where the prediction error $E_{ij} = q(\mathbf{e}_i, \mathbf{e}_j; \cdot) - R_{ij}$. If the prediction error is assumed to be depicted by a Gaussian distribution $p(E_{ij}) = \mathcal{N}(E_{ij}; 0, \sigma_e^2)$ for every visited joint action (i, j) , then the learning objective for the Q function can be formulated as:

$$\begin{aligned} & \mathcal{J}^{(T)}(\mathbf{W}, \mathbf{b}, \mathbf{c}, d) \\ &= \mathbb{E}_{\{(i^{(t)}, j^{(t)}) \sim p_e^{(t)}\}_{t=1}^T} \log \left(p(\mathbf{W}) \prod_{t=1}^T p(E_{i^{(t)} j^{(t)}}) \right) \\ &= \sum_{t=1}^T \mathbb{E}_{(i, j) \sim p_e^{(t)}} [\log \mathcal{N}(q(\mathbf{e}_i, \mathbf{e}_j; \mathbf{W}, \mathbf{b}, \mathbf{c}, d) - R_{ij}; 0, \sigma_e^2)] \\ & \quad + \log \mathcal{N}(\mathbf{W}; 0, \sigma_w^2 I) + \text{Const.} \end{aligned} \quad (4)$$

We use $q_{\mathcal{J}^{(T)}}(\mathbf{W}, \mathbf{b}, \mathbf{c}, d)$ to denote the learned joint Q function that maximizes $\mathcal{J}^{(T)}$ at step T . $q_{\mathcal{J}^{(T)}}(\mathbf{W}, \mathbf{b}, \mathbf{c}, d)$ is determined by the exploration policy $p_e^{(t)}$ and the exploration steps T . Then we have the following theorem for the uniform exploration strategy.

Theorem 4.2 (uniform exploration). *Assume $\delta \leq \frac{1}{6}$, $U \geq 3$. Using a uniform exploration policy in the climb game $G_f(2, 0, U)$, it can be proved that $q_{\mathcal{J}^{(T)}}(\mathbf{W}, \mathbf{b}, \mathbf{c}, d)$ will become equivalently optimal only after $T = \Omega(|\mathcal{A}| \delta^{-1})$ steps. When $\delta = 1$, $T = O(1)$ steps suffice to learn the equivalently optimal joint Q function, suggesting the inefficiency of uniform exploration is due to a large set of sub-optimal NEs.*

The intuition behind Theorem 4.2 is that the hardness of exploration in climb games largely comes from the sparsity of solutions: a set of sub-optimal NEs exist but there is only a single Pareto optimal NE. Learning the joint Q function can be influenced by the sub-optimal NEs. And if the exploration attempts are not well coordinated, a lot of zero reward would be encountered, making it hard to find the Pareto optimal NE. We also remark that uniform exploration can be particularly inefficient since the term $|\mathcal{A}|$ can be exponentially large in a multi-agent system. This indicates that more efficient exploration can potentially be achieved by reducing the search space and identifying a smaller "critical" subspace.

To formally prove Theorem 4.2, we define f_1, f_2, f_3 as the step-averaged probability of taking the joint action in optimal NE, sub-optimal NE and zero-reward, respectively. We show that to make

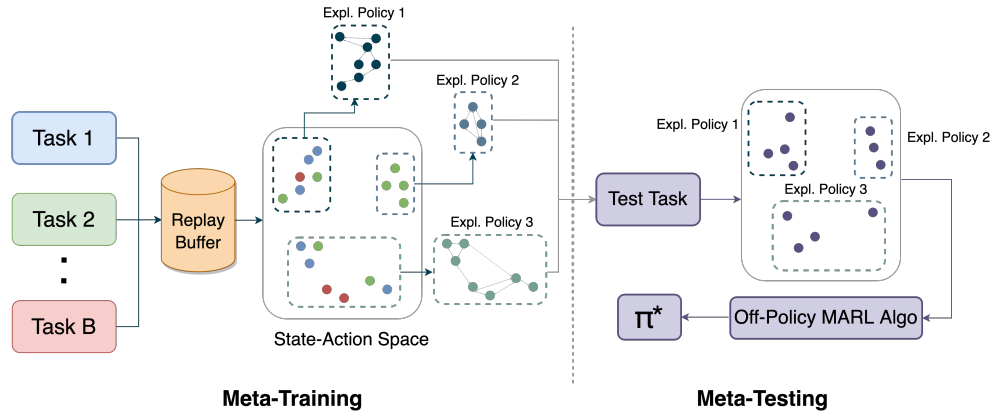


Figure 2: MESA’s meta-learning framework. In the meta-training stage, MESA learns exploration policies to cover the high-rewarding subspace. In the meta-testing stage, MESA uses the learned exploration policies to assist the learning in an unseen task. Each color corresponds to a different task, and the colored points represent the high-rewarding joint state-action pairs collected in that task.

the joint Q function equivalently optimal, there is a necessary condition that f_1, f_2, f_3 should follow. When T is not large enough, this condition cannot be satisfied. Detailed proof is in Appendix A.2.

Next, we consider the case of another popular exploration paradigm, ϵ -greedy exploration.

Theorem 4.3 (ϵ -greedy exploration). *Assume $\delta \leq \frac{1}{32}, U \geq 4, U \geq \sigma_w \sigma_e^{-1}$. In the climb game $G_f(2, 0, U)$, under ϵ -greedy exploration with fixed $\epsilon \leq \frac{1}{2}, q_{\mathcal{G}(T)}(\mathbf{W}, \mathbf{b}, \mathbf{c}, d)$ will become equivalently optimal only after $T = \Omega(|\mathcal{A}|\delta^{-1}\epsilon^{-1})$ steps. If $\epsilon(t) = 1/t$, it requires $T = \exp(\Omega(|\mathcal{A}|\delta^{-1}))$ exploration steps to be equivalently optimal.*

The proof is similar to that of Theorem 4.2 (detailed in Appendix A.3). By comparing 4.2 and 4.3, ϵ -greedy results in even poorer exploration efficiency than uniform exploration. Note the ϵ -greedy strategy is training policy specific, i.e., the exploration behavior varies as the training policy changes. Theorem 4.3 suggests that when the policy is sub-optimal, the induced ϵ -greedy exploration strategy can be even worse than uniform exploration. Hence, it can be beneficial to adopt a separate exploration independent from the training policy.

The above analysis shows that common exploration strategies like uniform exploration or ϵ -greedy exploration are inefficient for such a simple game and the main reason is that it requires coordination between different agents to reach high-rewarding states, but naive exploration strategies lack such cooperation.

4.2 Structured Exploration

We will show that it is possible to design a better exploration strategy with some prior knowledge of the climb game structure. Consider a specific structured exploration strategy $p_e^{(t)}(i, j) = U^{-1} \mathbb{1}_{i=j}$, where both agents always choose the same action. With such a strategy, we can quickly find the optimal solution to the game. More formally, we have the following theorem.

Theorem 4.4 (structured exploration). *In the climb game $G_f(2, 0, U)$, under structured exploration $p_e^{(t)}(i, j) = U^{-1} \mathbb{1}_{i=j}, q_{\mathcal{G}(T)}(\mathbf{W}, \mathbf{b}, \mathbf{c}, d)$ is equivalently optimal at step $T = O(1)$.*

Theorem 4.4 shows the efficiency of exploration can be greatly improved if the exploration strategy captures a proper structure of the problem, i.e., all agents taking the same action. We further remark that by considering a set of similar climb games $\mathcal{G} = \{G_f(2, u, U)\}_{u=0}^{U-1}$, the structured exploration strategy $p_e^{(t)}(i, j) = U^{-1} \mathbb{1}_{i=j}$ can be interpreted as a uniform distribution over the optimal policies of this game set \mathcal{G} . This interesting fact suggests that we can first collect a set of similarly structured games and then derive effective exploration strategies from these similar games. Once a set of structured exploration strategies are collected, we can further adopt them for fast learning in a novel game with a similar problem structure. We take the inspiration here and develop a general meta-exploration algorithm in the next section.

5 METHOD

We detail our method *Cooperative Meta-Exploration in Multi-Agent Learning through Exploiting State-Action Space Structure* (MESA) for cooperative multi-agent learning. As shown in Figure 2, MESA consists of a meta-training stage (Algo. 1) and a meta-testing stage (Algo. 2). In the meta-training stage, MESA learns exploration policies by training in a batch of training tasks that share intrinsic structures in the state-action space. In the meta-testing stage, MESA utilizes the meta-learned exploration policies to assist learning in an unseen task sampled from the distribution of the training tasks.

5.1 Meta-Training

The meta-training stage contains two steps: 1) identify the high-rewarding state-action subspace, and 2) train a set of exploration policies using the subspace-induced rewards.

5.1.1 Identifying High-Rewarding Joint State-Action Subspace. For each training task \mathcal{T}_i , we collect experiences $\mathcal{D}_i = \{(s_t, \mathbf{a}_t, r_t, s_{t+1})\}$.

Algorithm 1 MESA: Meta-Training

Input: Meta-training tasks $\{\mathcal{T}_i\}_{i=1}^B \sim p(\mathcal{T})$, off-policy MARL algorithm f , distance metric $\|\cdot\|_{\mathcal{F}}$
Parameter: #policies E , threshold R^* , horizon h
Output: Exploration policies $\{\pi_e^i\}_{i=1}^E$

- 1: $\mathcal{M}_* \leftarrow \emptyset$, global pseudo-count $\hat{N} \leftarrow 0$
- 2: **for** $i = 1$ to B **do**
- 3: Initialize policy π_θ
- 4: Train π_θ with f and collect dataset $D_i = \{(s_t, \mathbf{a}_t, r_t, s_{t+1})\}$
- 5: $\mathcal{M}_* \leftarrow \mathcal{M}_* \cup \{\tau \mid R(\tau) \geq R^*, \tau \in D_i\}$
- 6: **end for**
- 7: **for** $i = 1$ to E **do**
- 8: Initialize exploration policy π_e^i
- 9: **while** π_e^i 's training not converged **do**
- 10: Initialize N as \hat{N} , $\mathcal{D} \leftarrow \emptyset$
- 11: **for** $t = 0$ to $h-1$ **do**
- 12: Execute $\mathbf{a}_t \sim \pi_e^i(s_t)$, and observe $(s_t, \mathbf{a}_t, r_t, s_{t+1})$
- 13: Calculate \hat{r}_t based on Eq. 5 or 6
- 14: Store $(s_t, \mathbf{a}_t, \hat{r}_t, s_{t+1})$ into \mathcal{D}
- 15: $N(\phi(s_t, \mathbf{a}_t)) \leftarrow N(\phi(s_t, \mathbf{a}_t)) + 1$
- 16: **end for**
- 17: Optimize policy π_e^i with algorithm f
- 18: **end while**
- 19: Update \hat{N} using \mathcal{D}
- 20: **end for**
- 21: **return** $\{\pi_e^i\}_{i=1}^E$

If the reward r_t is higher than a threshold R^* , we call this joint state-action pair (s_t, \mathbf{a}_t) valuable and store it into a dataset \mathcal{M}_* . For goal-oriented tasks where $r = \mathbb{1}_{s=goal}$, the threshold can be set as $R^* = 1$. For other tasks, the threshold can be set as a hyperparameter, for example, a certain percentile of all collected rewards. A smaller R^* results in a larger identified subspace but a less efficient exploration policy.

The data stored in \mathcal{M}_* is highly diversified since it comes from all the B training tasks, which are expected to share an intrinsic structure. We expect that with this intrinsic structure, the high-rewarding joint state-action pairs fall into some low-dimensional subspace. In the simplest case, they may form several dense clusters, or many of them lie in a hyperplane. Even if the subspace is not easily interpretable to humans, it may still be effectively ‘‘covered’’ by a set of exploration policies (to be found in the subsequent step).

We also explicitly deal with the reward sparsity problem by assigning a positive reward to a joint state-action pair (s_t, \mathbf{a}_t) if it has zero reward but leads to a valuable state-action pair $(s_{t'}, \mathbf{a}_{t'})$ later in the same trajectory. We also put these relabeled pairs into the dataset \mathcal{M}_* . Let $t' = \arg \min_{t' > t} [r_{t'} > 0]$, we therefore have the following densified reward function

$$\hat{r}_t = \begin{cases} \gamma^{t'-t} \cdot r_{t'}, & r_t = 0, \\ r_t, & r_t > 0. \end{cases} \quad (5)$$

Algorithm 2 MESA: Meta-Testing

Input: Test task $\hat{\mathcal{T}}$, meta-trained exploration policies $\{\pi_e^i\}_{i=1}^E$, off-policy MARL algorithm f
Parameter: horizon h
Output: Policy π_θ for task $\hat{\mathcal{T}}$

- 1: Initialize policy π_θ , $\mathcal{D} = \emptyset$, annealing ϵ
- 2: **while** not converged **do**
- 3: Determine p_e under annealing probability schedule ϵ
- 4: Choose policy to perform rollouts by

$$\pi_d = \begin{cases} \pi_e \sim \mathcal{U}(\{\pi_e^i\}_{i=1}^E), & \text{w.p. } p_e \\ \pi_\theta, & \text{otherwise.} \end{cases}$$
- 5: **for** $t = 0$ to $h-1$ **do**
- 6: Execute $\mathbf{a}_t \sim \pi_d(s_t)$.
- 7: Observe transition $(s_t, \mathbf{a}_t, r_t, s_{t+1})$.
- 8: $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, \mathbf{a}_t, r_t, s_{t+1})$
- 9: **end for**
- 10: Optimize π_θ with algorithm f on replay buffer \mathcal{D}
- 11: **end while**
- 12: **return** π_θ

5.1.2 Learning Exploration Policies. In this step, we aim to learn a diverse set of exploration policies to cover the identified high-rewarding joint state-action subspace. We use a distance metric $\|\cdot\|_{\mathcal{F}}$ (e.g., l_2 distance) to determine whether two state-action pairs are close. Then if a visited joint state-action pair (s, \mathbf{a}) is close enough to the identified subspace \mathcal{M}_* , i.e., $\min_{d \in \mathcal{M}_*} \|(s, \mathbf{a}), d\|_{\mathcal{F}} < \epsilon$, it would be assigned a derived positive reward \hat{r} . Increasing the value of B in the collection step would generally result in a more accurate distance measurement. However, this comes at the cost of making the minimization calculation more computationally expensive.

To encourage a broader coverage of the subspace and to avoid mode collapse, the reward assignment scheme ensures that repeated visits to similar joint state-action pairs within one trajectory would result in a decreasing reward for each visit. Similar to [37], we adopt a pseudo-count function N with a hash function $\phi(s, \mathbf{a})$ to generalize between similar joint state-action pairs. We then apply a decreasing function $f_d : \mathcal{N} \mapsto [0, 1]$ on the trajectory-level pseudo-count $N(\phi((s, \mathbf{a})))$. The resulted reward assignment scheme is defined as follows:

$$\tilde{r}_t = \hat{r}_t f_d(N(\phi((s_t, \mathbf{a}_t)))) \left[\mathbb{1}_{\min_{d \in \mathcal{M}_*} \|(s_t, \mathbf{a}_t), d\|_{\mathcal{F}} < \epsilon} \right] \quad (6)$$

After one exploration policy is trained with this reward, we will train a new policy to cover the part of the identified subspace that has not yet been covered. This is achieved by having a global pseudo-count \hat{N} which is updated after training each exploration policy using its visitation counts and is maintained throughout the training of all exploration policies. This iterative process continues until the subspace is well-covered by the set of trained exploration policies.

5.2 Meta-Testing

During meta-testing, MESA uses the meta-learned exploration policies $\{\pi_e^i\}_{i=1}^E$ to assist the training of any generic off-policy MARL algorithm on a test-time task $\hat{\mathcal{T}}$. Specifically, for each rollout episode, we choose with probability ϵ to execute one uniformly sampled exploration policy $\pi_e \sim \mathcal{U}(\{\pi_e^i\}_{i=1}^E)$. For the best empirical performance, we also adopt an annealing schedule $\epsilon : T \mapsto [0, 1]$ so that the exploration policies provide more rollouts at the initial stage of the training and are gradually turned off later.

Here we further provide some analysis of deploying the meta-learned exploration policy on unseen testing tasks.

Theorem 5.1 (Exploration during Meta-Testing). *Consider goal-oriented tasks with goal space $\mathcal{G} \subseteq \mathcal{S}$. Assume the training and testing goals are sampled from the distribution $p(x)$ on \mathcal{G} , and the dataset has N i.i.d. goals sampled from a distribution $q(x)$ on \mathcal{S} . If the exploration policy generalizes to explore ϵ nearby goals for every training sample, we have that the testing goal is not explored with probability at most*

$$P_{fail} \approx \int p(x)(1 - \epsilon q(x))^N dx \leq O\left(\frac{KL(p||q) + \mathcal{H}(p)}{\log(\epsilon N)}\right). \quad (7)$$

Theorem 5.1 shows that the good performance of meta-learned exploration policy relies on 1) a small difference between the training and testing distribution; and 2) a structured, e.g., low-dimensional, high-rewarding subspace \mathcal{G} to reduce $\mathcal{H}(p)$. And when uniformly sampling the training data, $KL(p||q)$ is bounded by $\log \Omega_{\mathcal{G}}$ in our method. This term, however, can be up to $\log \Omega_{\mathcal{S}}$ with an uncoordinated exploration on the joint state space \mathcal{S} , where $\Omega_{\mathcal{S}}$ can be exponentially larger than $\Omega_{\mathcal{G}}$.

5.3 Implementation Detail of MESA

We choose MADDPG, following the centralized training with decentralized execution (CTDE) paradigm, as the off-policy MARL algorithm for MESA since it can be applied to both discrete and continuous action space, as shown in its original paper [23]. We use a clustering mapping f_c as the hash function ϕ so that the dataset \mathcal{M}_* is clustered into C clusters defined by the clustering function $f_c : \mathcal{S} \times \mathcal{A} \mapsto [C]$. The cluster mapping is implemented with the KMeans clustering algorithm [22]. The number of exploration policies to learn is viewed as a hyperparameter. See the Appendix for detailed hyperparameter settings.

6 EXPERIMENTS

Our experimental evaluation aims to answer the following questions: (1) Are the meta-learned exploration policies capable of achieving more efficient exploration during meta-testing on newly sampled tasks in matrix climb game variants (Section 6.2) and high-dimensional domains (Section 6.3 and 6.4)? (2) Can these meta-learned exploration policies successfully generalize to unseen test-time tasks from a more challenging (e.g., with more agents) test task distribution which is different the training task distribution (Section 6.5)?

6.1 Evaluation Setup

Compared Methods. We compare to 3 multi-agent reinforcement learning algorithms: MADDPG [23], MAPPO [41], and QMIX [33],

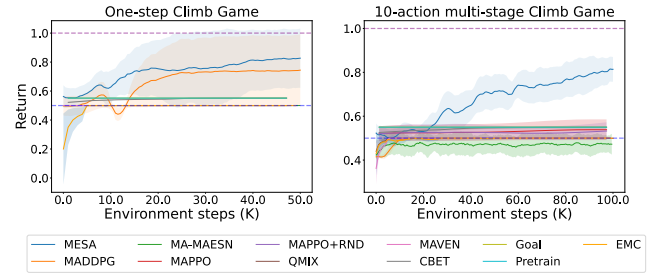


Figure 3: Learning curve of the two climb game variants w.r.t number of environment steps. The return is averaged over timesteps for the multi-stage games. The dotted lines indicate the suboptimal return of 0.5 (purple) and the optimal return 1 (blue) for each agent.

to measure the effectiveness of our exploration policies. We also compare to 3 multi-agent exploration algorithm: MAVEN [24], MAPPO with RND exploration [5], and EMC [43]. To compare with baselines that adopt a similar meta-training stage, we add two naive meta-learning baselines, including one with an unconditioned shared policy, which is trained over all training tasks, and one with a goal-conditioned policy, which takes the target landmarks as parts of the input. We also adapt the single-agent meta-RL algorithm MAESN [14] to the multi-agent setting. Finally, we adapt the single-agent C-BET [26] to multi-agent settings based on MAPPO. The training and testing tasks are as defined in Section 6.1. Please refer to the Appendix for more visualization and experimental results.

Environments. We experiment on the Climb Game, Multi-agent Particle Environment (MPE) [23], and multi-agent MuJoCo [29], on which generating a distribution of meta-training tasks $p(\mathcal{T})$ is feasible.

6.2 Climb Game Variants

First, we consider task spaces consisting of variants of the aforementioned climb games. We extend previous climb game to (1) **one-step climb game** $G(n, k, u, U)$, which is a n -player game with U actions for each player, and the joint reward is 1 if $\#u = k$, $1 - \delta$ if $\#u = 0$, and 0 otherwise. The task space $\mathcal{T}_U^{\text{one}}$ consists of all one-step climb games that contain two players and U actions; (2) **multi-stage climb game**, which is an S -stage game where each stage is a one-stage climb game with the same number of available actions. Each stage t has its own configuration (k_t, u_t) of the one-stage climb game $G(2, k_t, u_t, U)$. Agents observe the history of joint actions and the current stage t . The task space $\mathcal{T}_{S,U}^{\text{multi}}$ consists of all multi-stage climb games with S stages and U actions. In our experiments, we use $\mathcal{T}_{10}^{\text{one}}$ and $\mathcal{T}_{5,10}^{\text{multi}}$ as the task space for the one-step and multi-stage Climb Games. We choose uniformly at random ten training tasks and three different test tasks from the task space \mathcal{T} , and we keep $\delta = \frac{1}{2}$ as in the classic climb games.

Results on Climb Game Variants. For the matrix games, we additionally compare with MA-MAESN, which is our adaptation of the original single-agent meta-learning algorithm MAESN [14] to the multi-agent scenario. In the single-step matrix game, MESA exhibits better performance, being able to find the optimal reward

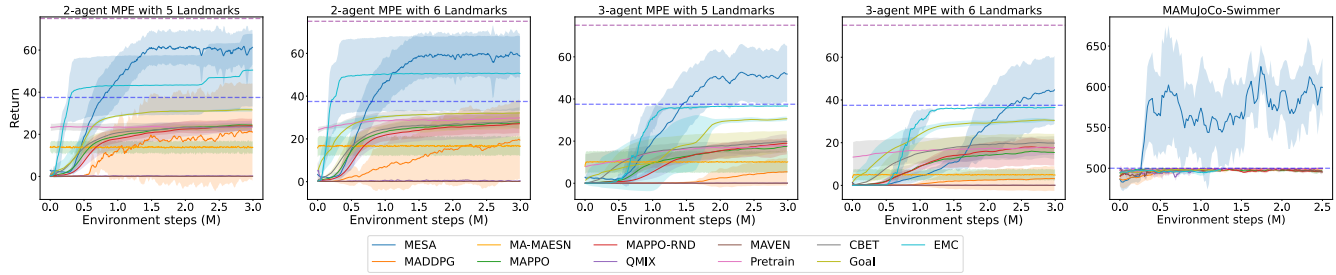


Figure 4: Learning curves of MESA and the compared baselines w.r.t the number of environment interactions during the meta-testing stage in the MPE domain and the multi-agent MuJoCo environment Swimmer. The two dotted lines indicate the ideal optimal (purple) and sub-optimal (blue) return summed over timesteps. A return above the blue line would typically indicate that the agents are able to learn the optimal strategy.

in some harder tasks when $k = 2$, while other baselines are stuck at the sub-optimal reward for almost all tasks.

In the more challenging 10-action multi-stage game where task space is exponentially larger, MESA outperforms all compared algorithms by a large margin. With the help of the exploration policies that have learned the high-rewarding joint action pairs, MESA quickly learns the optimal joint action for each stage and avoids being stuck at the sub-optimal.




Figure 6: Visualization of structured exploration behaviors discovered by the meta-trained exploration policy in MESA.

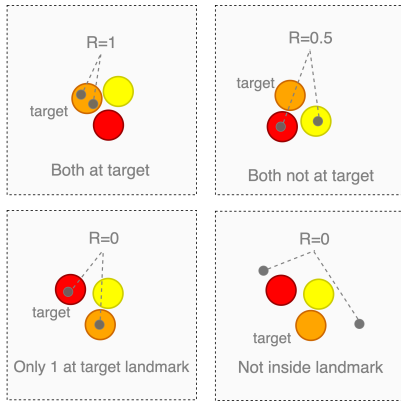


Figure 5: Visualizations of a 2-player 3-landmark MPE climb game.

6.3 MPE Domain

We extend the matrix climb games to MPE [23], which has a continuous high-dimensional state space. Agents must first learn to reach the landmarks under sparse rewards and then learn to play the climb games optimally.

In a MPE Climb Game $\tilde{G}(n, k, u, U, \{L_j\}_0^{U-1})$ (Figure 5), there are U non-overlapping landmarks with positions $\{L_j\}_{j=0}^{U-1}$. The reward is non-zero only when every agent is on some landmark. Agents will be given a reward of 1 if there are exactly k agents located on the u -th landmark (target landmark), and a suboptimal reward of $1 - \delta$ will be given when none of the agents are located on the target landmark. Otherwise, the reward will be zero. As before, u and k are not present in the observation and can only be inferred

from the received reward. A task space $\mathcal{T}_{n,U}^{MPE}$ consists of all MPE climb games with n players and U landmarks. We evaluate MESA on the 2-agent tasks ($\mathcal{T}_{2,5}^{MPE}$ and $\mathcal{T}_{2,6}^{MPE}$) and 3-agent tasks ($\mathcal{T}_{3,5}^{MPE}$ and $\mathcal{T}_{3,6}^{MPE}$) while fixing $k = 2$. Each sampled training and testing task has a different configuration of landmark positions.

Adaptation Performance in MPE. We show in Figure 4 the learning curve of our approach MESA compared with the aforementioned baseline methods. MESA outperforms the compared baselines by a large margin, being able to coordinately reach the task landmark quickly, as evidenced by the near-optimal reward. Even when combined with RND-based exploration, MAPPO easily sticks to the sub-optimal equilibrium. Value-based methods like QMIX and MAVEN are unable to learn the correct Q -function because the reward is quite sparse before agents can consistently move themselves to a landmark. EMC sometimes jumps out of the suboptimal equilibrium with curiosity-driven exploration, but the performance is not robust. Furthermore, as the meta-learning baselines only learn the sub-optimal behavior during meta-training, they fail to learn the optimal equilibrium during test time and quickly converge to the suboptimal equilibrium.

Visualization of Exploration Policies. To answer question (2), we visualize the learned exploration policies in a 2-agent 3-landmark MPE task in Figure 6. We can see that the learned exploration policy consecutively visited the 3 landmarks within 20 timesteps in one trajectory.

6.4 Multi-agent MuJoCo Environments

We also extend the matrix climb games to multi-agent MuJoCo environments [29]. We consider specifically the 2-agent Swimmer environment where each agent is a hinge on the swimmer’s body,

and each agent’s action is the amount of torque applied to hinge rotors. The extension considers the angles between the two hinges and the body segments. Each task in the task space is a target angle such that a reward of 1 will be given only if the two angles are both close to the target angles, a 0.5 suboptimal reward is given if none of two angles are close to the target, and a reward of 0 if only one of the two angles are close.

This multi-agent environment is extremely hard as agents are very likely to converge to the suboptimal reward of 0.5, which is confirmed by the results that none of the baselines were able to find the optimal equilibrium in Figure 4. Therefore, MESA vastly outperforms all the compared baselines by learning a final policy that frequently reaches the target angle.

6.5 Generalization Performance of MESA

In this section, our goal is to evaluate the generalization performance of the meta-trained exploration policy in scenarios where the meta-training and meta-testing task distributions are different. In particular, we focus on the setting where the test-time tasks are *more challenging* than the training-time tasks and examine how an exploration policy learned from simpler tasks can boost training performances on harder tasks.

The test task here is uniform on the 3-agent high-difficulty MPE Climb games. The task difficulty is defined by the average pairwise distances between the landmark positions and the initial positions of the agents. We consider two simpler training task distributions, including (1) a 2-agent setting with the same difficulty, and (2) a 3-agent setting with a lower difficulty. In both settings, the meta-training tasks are less challenging than the test-time tasks. For evaluation, the meta-trained exploration policy from each setting will be directly applied to assist the training on the more challenging test-time tasks, *without any fine-tuning*.

We modified the neural network architecture by adopting an attention layer in both actor and critic to ensure they are compatible with a varying number of agents. The attention mechanism acts as an aggregation function between the relative positions of the other agents and its own relative position to the landmarks to handle the varying observation dimensions. Additionally, we employed behavior cloning (BC) [30] on the rollouts of the exploration policies as a warm-up to accelerate learning of the final policy.

In Figure 7, we present the generalization results from our study. We evaluate the zero-shot generalization ability of the meta-exploration policy by measuring the average number of high-reward transitions hit in a test task randomly sampled from the test task distribution. As shown on the left of Figure 7, the meta-exploration policies are able to explore the test-time tasks much more efficiently than a random exploration policy, even on test-time tasks that are drawn from a harder task distribution. Notably, the generalization ability increases with the number of exploration policies (B). Using the meta-exploration policies trained on the simpler tasks, MESA is able to consistently reach the high-reward region in the unseen hard 3-agent tasks, as opposed to the vanilla MADDPG algorithm that only learns the sub-optimal equilibrium. We also see that with an increasing number of meta-exploration policies, the performance of MESA increases, but the improvement becomes marginal, while the meta-training time increases linearly with E .

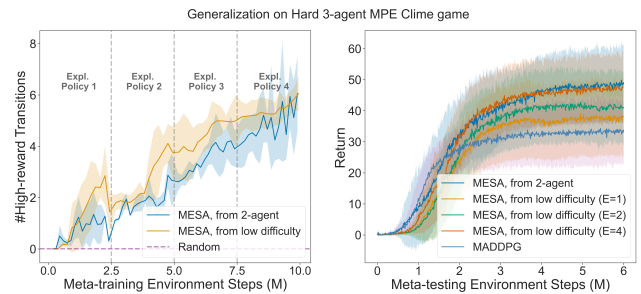


Figure 7: Generalization results of MESA on the hard 3-agent MPE Climb game. Left: Zero-shot generalizability of the meta-exploration policies, measured by the number of visitations on high-reward transitions per episode on the test tasks. The purple dotted line corresponds to the random exploration policy. The plot shows the concatenated training curves for all exploration policies. Right: Learning curves of MESA under different settings using the meta-exploration policies trained on the two different training-task distributions.

7 CONCLUSIONS

This paper introduces a meta-exploration method, MESA, for multi-agent learning. The key idea is to learn a diverse set of exploration policies to cover the high-rewarding state-action subspace and achieve efficient exploration in an unseen task. MESA can work with any off-policy MARL algorithm, and empirical results confirm the effectiveness of MESA in climb games, MPE environments, and multi-agent MuJoCo environments and its generalizability to more complex test-time tasks.

ACKNOWLEDGMENTS

This research is supported in part by NSF IIS-2046640 (CAREER) and Sloan Research Fellowship. We thank NVIDIA for providing computing resources. Zhicheng Zhang is supported in part by SCS Dean’s Fellowship. The funders have no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

REFERENCES

- [1] Jacob Andreas, Dan Klein, and Sergey Levine. 2017. Modular multitask reinforcement learning with policy sketches. In *ICML*. PMLR, 166–175.
- [2] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. *arXiv preprint arXiv:1707.01495* (2017).
- [3] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. Unifying count-based exploration and intrinsic motivation. *NeurIPS* 29 (2016), 1471–1479.
- [4] Avrim Blum and Ronald Rivest. 1988. Training a 3-node neural network is NP-complete. *NeurIPS* 1 (1988).
- [5] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. 2018. Exploration by random network distillation. In *ICLR*.
- [6] Cédric Colas, Olivier Sigaud, and Pierre-Yves Oudeyer. 2017. GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms. In *ICML*.
- [7] Ron Dorfman, Idan Shenfeld, and Aviv Tamar. 2020. Offline Meta Learning of Exploration. *arXiv preprint arXiv:2008.02598* (2020).
- [8] Simon Du and Jason Lee. 2018. On the power of over-parametrization in neural networks with quadratic activation. In *ICML*. PMLR, 1329–1338.
- [9] Yan Duan, John Schulman, Xi Chen, Peter L. Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779* (2016).

- [10] Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. 2021. First return, then explore. *Nature* 590, 7847 (2021), 580–586.
- [11] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. 2018. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *ICML*. PMLR, 1407–1416.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the 34th ICML (PMLR, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 1126–1135. <https://proceedings.mlr.press/v70/finn17a.html>
- [13] Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. 2016. Continuous deep q-learning with model-based acceleration. In *ICML*. PMLR, 2829–2838.
- [14] Abhishek Gupta, Russell Mendonca, Yuxuan Liu, Pieter Abbeel, and Sergey Levine. 2018. Meta-Reinforcement Learning of Structured Exploration Strategies. *NIPS 2018* (2018), 5302–5311.
- [15] Tarun Gupta, Anuj Mahajan, Bei Peng, Wendelin Böhmer, and Shimon Whiteson. 2021. Uneven: Universal value exploration for multi-agent reinforcement learning. In *ICML*. PMLR, 3930–3941.
- [16] Matteo Hessel, Hubert Soyer, Lasse Espeholt, Wojciech Czarnecki, Simon Schmitt, and Hado van Hasselt. 2019. Multi-task deep reinforcement learning with popart. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 3796–3803.
- [17] Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Duenez-Guzman, Antonio Garcia Castaneda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. *NIPS 2018* 31 (2018).
- [18] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, DJ Strouse, Joel Z Leibo, and Nando De Freitas. 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *ICML*. PMLR, 3040–3049.
- [19] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro A Ortega, DJ Strouse, Joel Z Leibo, and Nando de Freitas. 2018. Intrinsic social motivation via causal influence in multi-agent RL. (2018).
- [20] Lin Lan, Zhenguo Li, Xiaohong Guan, and Pinghui Wang. 2019. Meta reinforcement learning with task embedding and shared policy. *arXiv preprint arXiv:1905.06527* (2019).
- [21] Iou-Jen Liu, Unnat Jain, Raymond A Yeh, and Alexander Schwing. 2021. Cooperative exploration for multi-agent deep reinforcement learning. In *ICML*. PMLR, 6826–6836.
- [22] Stuart Lloyd. 1982. Least squares quantization in PCM. *IEEE transactions on information theory* 28, 2 (1982), 129–137.
- [23] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *NIPS*.
- [24] Anuj Mahajan, Tabish Rashid, Mikayel Samvelyan, and Shimon Whiteson. 2019. MAVEN: Multi-Agent Variational Exploration. In *NeurIPS*, Vol. 32. 7613–7624.
- [25] Georg Ostrovski, Marc G Bellemare, Aaron Oord, and Rémi Munos. 2017. Count-based exploration with neural density models. In *ICML*. PMLR, 2721–2730.
- [26] Simone Parisi, Victoria Dean, Deepak Pathak, and Abhinav Gupta. 2021. Interesting object, curious agent: Learning task-agnostic exploration. *NeurIPS* 34 (2021), 20516–20530.
- [27] Emilio Parisotto, Lei Jimmy Ba, and Ruslan Salakhutdinov. 2016. Actor-Mimic: Deep Multitask and Transfer Reinforcement Learning. In *ICLR (Poster)*.
- [28] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *ICML*. PMLR, 2778–2787.
- [29] Bei Peng, Tabish Rashid, Christian Schroeder de Witt, Pierre-Alexandre Kamienny, Philip Torr, Wendelin Böhmer, and Shimon Whiteson. 2021. Facmac: Factored multi-agent centralised policy gradients. *NeurIPS* 34 (2021).
- [30] Dean A Pomerleau. 1991. Efficient training of artificial neural networks for autonomous navigation. *Neural computation* 3, 1 (1991), 88–97.
- [31] Vitchyr H Pong, Ashvin Nair, Laura Smith, Catherine Huang, and Sergey Levine. 2021. Offline Meta-Reinforcement Learning with Online Self-Supervision. *arXiv preprint arXiv:2107.03974* (2021).
- [32] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. 2019. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *ICML*. PMLR, 5331–5340.
- [33] 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*. PMLR, 4295–4304.
- [34] Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degraeve, Tom Wiele, Vlad Mnih, Nicolas Heess, and Jost Tobias Springenberg. 2018. Learning by playing solving sparse reward tasks from scratch. In *ICML*. PMLR, 4344–4353.
- [35] Lukas Schäfer, Filippos Christianos, Josiah P Hanna, and Stefano V Albrecht. 2022. Decoupled Reinforcement Learning to Stabilise Intrinsically-Motivated Exploration. In *Proceedings of the 21st AAMAS*. 1146–1154.
- [36] Avi Singh, Huihan Liu, Gaoyue Zhou, Albert Yu, Nicholas Rhinehart, and Sergey Levine. 2020. Parrot: Data-Driven Behavioral Priors for Reinforcement Learning. In *ICLR*.
- [37] Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. 2017. # exploration: A study of count-based exploration for deep reinforcement learning. In *31st NIPS*, Vol. 30. 1–18.
- [38] Pin Wang, Hanhan Li, and Ching-Yao Chan. 2019. Quadratic Q-network for learning continuous control for autonomous vehicles. *arXiv preprint arXiv:1912.00074* (2019).
- [39] Tonghan Wang, Jianhao Wang, Yi Wu, and Chongjie Zhang. 2019. Influence-Based Multi-Agent Exploration. In *ICLR*.
- [40] Tianbing Xu, Qiang Liu, Liang Zhao, and Jian Peng. 2018. Learning to explore with meta-policy gradient. *arXiv preprint arXiv:1803.05044* (2018).
- [41] Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu. 2021. The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games. [arXiv:2103.01955 \[cs.LG\]](https://arxiv.org/abs/2103.01955)
- [42] Jin Zhang, Jianhao Wang, Hao Hu, Tong Chen, Yingfeng Chen, Changjie Fan, and Chongjie Zhang. 2021. Metacure: Meta reinforcement learning with empowerment-driven exploration. In *ICML*. PMLR, 12600–12610.
- [43] 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. *NeurIPS* 34 (2021).
- [44] Luisa Zintgraf, Kyriacos Siarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. 2020. VariBAD: A Very Good Method for Bayes-Adaptive Deep RL via Meta-Learning. In *ICLR*.