
Cooperative Exploration for Multi-Agent Deep Reinforcement Learning

Iou-Jen Liu¹ Unnat Jain¹ Raymond A. Yeh¹ Alexander G. Schwing¹

Abstract

Exploration is critical for good results in deep reinforcement learning and has attracted much attention. However, existing *multi-agent* deep reinforcement learning algorithms still use mostly noise-based techniques. Very recently, exploration methods that consider cooperation among multiple agents have been developed. However, existing methods suffer from a common challenge: agents struggle to identify states that are worth exploring, and hardly coordinate exploration efforts toward those states. To address this shortcoming, in this paper, we propose cooperative multi-agent exploration (CMAE): agents share a common goal while exploring. The goal is selected from multiple projected state spaces via a normalized entropy-based technique. Then, agents are trained to reach this goal in a coordinated manner. We demonstrate that CMAE consistently outperforms baselines on various tasks, including a sparse-reward version of the multiple-particle environment (MPE) and the Starcraft multi-agent challenge (SMAC).

1. Introduction

Multi-agent reinforcement learning (MARL) is an increasingly important field. Indeed, many real-world problems are naturally modeled using MARL techniques. For instance, tasks from areas as diverse as robot fleet coordination (Swamy et al., 2020; Hüttenrauch et al., 2019) and autonomous traffic control (Bazzan, 2008; Sunehag et al., 2018) fit MARL formulations.

To address MARL problems, early work followed independent single-agent reinforcement learning work (Tampuu et al., 2015; Tan, 1993; Matignon et al., 2012). However, more recently, specifically tailored techniques such as monotonic value function factorization (QMIX) (Rashid et al., 2018), multi-agent deep deterministic policy gradient

(MADDPG) (Lowe et al., 2017), and counterfactual multi-agent policy gradients (COMA) (Foerster et al., 2018) have been developed. Those methods excel in a multi-agent setting because they address the non-stationary issue of MARL via a centralized critic. Despite those advances and the resulting reported performance improvements, a common issue remains: all of the aforementioned methods use exploration techniques from classical algorithms. Specifically, these methods employ noise-based exploration, *i.e.*, the exploration policy is a noisy version of the actor policy (Rashid et al., 2020a; Lowe et al., 2017; Foerster et al., 2016; Rashid et al., 2018; Yang et al., 2018).

It was recently recognized that use of classical exploration techniques is sub-optimal in a multi-agent reinforcement learning setting. Specifically, Mahajan et al. (2019) show that QMIX with ϵ -greedy exploration results in slow exploration and sub-optimality. Mahajan et al. (2019) improve exploration by conditioning an agent’s behavior on a shared latent variable controlled by a hierarchical policy. Even more recently, Wang et al. (2020) encourage coordinated exploration by considering the influence of one agent’s behavior on other agents’ behaviors.

While all of the aforementioned exploration techniques for multi-agent reinforcement learning significantly improve results, they suffer from two common challenges: (1) agents struggle to identify states that are worth exploring. Identifying under-explored states is particularly challenging when the number of agents increase, since the state and action space grows exponentially with the number of agents. (2) Agents don’t coordinate their exploration efforts toward under-explored states. To give an example, consider a Push-Box task, where two agents need to jointly push a heavy box to a specific location before observing a reward. In this situation, instead of exploring the environment independently, agents need to coordinate pushing the box within the environment to find the specific location.

To address both challenges, we propose cooperative multi-agent exploration (CMAE). To identify states that are worth exploring, we observe that, while the state space grows exponentially, the reward function typically depends on a small subset of the state space. For instance, in the aforementioned Push-Box task, the state space contains the location of agents and the box while the reward function only

¹University of Illinois at Urbana-Champaign, IL, U.S.A.. Correspondence to: Iou-Jen Liu <iliu3@illinois.edu>.

depends on the location of the box. To solve the task, exploring the box’s location is much more efficient than exploring the full state space. To encode this inductive bias into CMAE, we propose a bottom-up exploration scheme. Specifically, we project the high-dimensional state space to low-dimensional spaces, which we refer to as restricted spaces. Then, we gradually explore restricted spaces from low- to high-dimensional. To ensure the agents coordinate their exploration efforts, we select goals from restricted spaces and train the exploration policies to reach the goal. Specifically, inspired by Andrychowicz et al. (2017), we reshape the rewards in the replay buffer such that a positive reward is given when the goal is reached.

To show that CMAE improves results, we evaluate the proposed approach on two multi-agent environment suites: a discrete version of the multiple-particle environment (MPE) (Lowe et al., 2017; Wang et al., 2020) and the Starcraft multi-agent challenge (SMAC) (Samvelyan et al., 2019). In both environments, we consider both dense-reward and sparse-reward settings. Sparse-reward settings are particularly challenging because agents need to coordinate their behavior for extended timesteps before receiving any non-zero reward. CMAE consistently outperforms the state-of-the-art baselines in sparse-reward tasks. For more, please see our project page: <https://ioujenliu.github.io/CMAE>.

2. Preliminaries

We first define the multi-agent Markov decision process (MDP) in Sec. 2.1 and introduce the multi-agent reinforcement learning setting in Sec. 2.2.

2.1. Multi-Agent Markov Decision Process

A cooperative multi-agent system is modeled as a multi-agent Markov decision process (MDP). An n -agent MDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{Z}, \mathcal{O}, n, \gamma, H)$. \mathcal{S} is the state space of the environment. \mathcal{A} is the action space of each agent. At each time step t , each agent’s target policy π_i , $i \in \{1, \dots, n\}$, selects an action $a_i^t \in \mathcal{A}$. All selected actions form a joint action $\mathbf{a}^t \in \mathcal{A}^n$. The transition function \mathcal{T} maps the current state s^t and the joint action \mathbf{a}^t to a distribution over the next state s^{t+1} , *i.e.*, $\mathcal{T} : \mathcal{S} \times \mathcal{A}^n \rightarrow \Delta(\mathcal{S})$. All agents receive a collective reward $r^t \in \mathbb{R}$ according to the reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A}^n \rightarrow \mathbb{R}$. The objective of all agents’ policies is to maximize the collective return $\sum_{t=0}^H \gamma^t r^t$, where $\gamma \in [0, 1]$ is the discount factor, H is the horizon, and r^t is the collective reward obtained at timestep t . Each agent i observes local observation $o_i^t \in \mathcal{Z}$ according to the observation function $\mathcal{O} : \mathcal{S} \rightarrow \mathcal{Z}$. Note, observations usually reveal partial information about the state. For instance, suppose the state contains the location of agents, while the local observation

of an agent may only contain the location of other agents within a limited distance. All agents’ local observations form a joint observation \mathbf{o}^t .

2.2. Multi-Agent Reinforcement Learning

In this paper, we follow the standard centralized training and decentralized execution (CTDE) paradigm (Lowe et al., 2017; Rashid et al., 2018; Foerster et al., 2018; Mahajan et al., 2019; Liu et al., 2019b): at training time, the learning algorithm has access to all agents’ local observations, actions, and the state. At execution time, *i.e.*, at test time, each individual agent only has access to its own local observation.

The proposed CMAE is applicable to off-policy MARL methods (*e.g.*, Rashid et al., 2018; Lowe et al., 2017; Sunehag et al., 2018; Matignon et al., 2012; Liu et al., 2019b). In off-policy MARL, exploration policies $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^n$ are responsible for collecting data from the environment. The data in the form of transition tuples $(s^t, \mathbf{o}^t, \mathbf{a}^t, s^{t+1}, \mathbf{o}^{t+1}, r^t)$ is stored in a replay memory \mathcal{D} , *i.e.*, $\mathcal{D} = \{(s^t, \mathbf{o}^t, \mathbf{a}^t, s^{t+1}, \mathbf{o}^{t+1}, r^t)\}_t$. The target policies $\boldsymbol{\pi} = \{\pi_i\}_{i=1}^n$ are trained using transition tuples from the replay memory.

3. Coordinated Multi-Agent Exploration (CMAE)

In the following we first present an overview of CMAE before we discuss the method more formally.

Overview: The goal is to train the target policies $\boldsymbol{\pi} = \{\pi_i\}_{i=1}^n$ of n agents to maximize the environment episode return. Classical off-policy algorithms (Lowe et al., 2017; Rashid et al., 2018; Mnih et al., 2013; 2015; Lillicrap et al., 2016) typically use a noisy version of the target policies $\boldsymbol{\pi}$ as exploration policies. In contrast, in CMAE, we decouple exploration policies and target policies. Specifically, target policies are trained to maximize the usual external episode return. Exploration policies $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^n$ are trained to reach shared goals, which are under-explored states, as the job of an exploration policy is to collect data from those under-explored states.

To train the exploration policies, shared goals are required. How to choose shared goals from a high-dimensional state space? As discussed in Sec. 1, while the state space grows exponentially with the number of agents, the reward function often only depends on a small subset of the state space. Concretely, consider an n -agent Push-Box game in a $L \times L$ grid. The size of its state space is $L^{2(n+1)}$ (n agents plus box in L^2 space). However, the reward function depends only on the location of the box, whose state space size is L^2 . Obviously, to solve the task, exploring the location of the box is much more efficient than uniformly exploring the full state space.

Algorithm 1: Training with Coordinated Multi-Agent Exploration (CMAE)

Init: space tree T_{space} , counters \mathbf{c}
Init: exploration policies $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^n$, target policies $\boldsymbol{\pi} = \{\pi_i\}_{i=1}^n$, replay buffer \mathcal{D}

- 1 **for** $episode = 1 \dots E$ **do**
- 2 Reset the environment. Observe state s^1 and observations $\mathbf{o}^1 = (o_1^1, \dots, o_n^1)$
- 3 **for** $t = 1 \dots H$ **do**
- 4 Select \mathbf{a}^t using a mixture of exploration and target policies $\alpha\boldsymbol{\mu} + (1 - \alpha)\boldsymbol{\pi}$. α decreases linearly to 0
- 5 $r^t, s^{t+1}, \mathbf{o}^{t+1} = \text{environment.step}(\mathbf{a}^t)$
- 6 Add transition tuple $\{s^t, \mathbf{o}^t, \mathbf{a}^t, s^{t+1}, \mathbf{o}^{t+1}, r^t\}$ to \mathcal{D}
- 7 UpdateCounters($\mathbf{c}, s^{t+1}, \mathbf{o}^{t+1}$)
- 8 TrainTarget($\boldsymbol{\pi}, \mathcal{D}$)
- 9 **if** $episode \bmod N = 0$ **then**
- 10 $g = \text{SelectRestrictedSpaceGoal}(\mathbf{c}, T_{\text{space}}, \mathcal{D}, \text{episode})$ \triangleright Select shared goal (Alg. 3)
- 11 TrainExp($\boldsymbol{\mu}, g, \mathcal{D}$) \triangleright Train exploration policies (Alg. 2)

Algorithm 2: Train Exploration Policies (Train-Exp)

input : exploration policies $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^n$, shared goal g , replay buffer \mathcal{D}

- 1 **for** $\{s^t, \mathbf{o}^t, \mathbf{a}^t, s^{t+1}, \mathbf{o}^{t+1}, r^t\} \in \mathcal{D}$ **do**
- 2 **if** s^t is the shared goal g **then**
- 3 $r^t = r^t + \hat{r}$
- 4 Update $\boldsymbol{\mu}$ using $\{s^t, \mathbf{o}^t, \mathbf{a}^t, s^{t+1}, \mathbf{o}^{t+1}, r^t\}$

To achieve this, CMAE first explores a low-dimensional restricted space \mathcal{S}_k of the state space \mathcal{S} , *i.e.*, $\mathcal{S}_k \subseteq \mathcal{S}$. Formally, given an M -dimensional state space \mathcal{S} , the restricted space \mathcal{S}_k associated with a set k is defined as

$$\mathcal{S}_k = \{\text{proj}_k(s) : \forall s \in \mathcal{S}\}, \quad (1)$$

where $\text{proj}_k(s) = (s_e)_{e \in k}$ ‘restricts’ the space to elements e in set k , *i.e.*, $e \in k$. Here, s_e is the e -th component of the full state s , and k is a set from the power set of $\{1, \dots, M\}$, *i.e.*, $k \in P(\{1, \dots, M\})$, where P denotes the power set. CMAE gradually moves from low-dimensional restricted spaces ($|k|$ small) to higher-dimensional restricted spaces ($|k|$ larger). This bottom-up space selection is formulated as a search on a space tree T_{space} , where each node represents a restricted space \mathcal{S}_k .

Alg. 1 summarizes this approach. At each step, a mixture of the exploration policies $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^n$ and target policies $\boldsymbol{\pi} = \{\pi_i\}_{i=1}^n$ is used to select actions (line 4). The resulting experience tuple is then stored in a replay buffer \mathcal{D} . Counters \mathbf{c} for each restricted space in the space tree T_{space} track how often a particular restricted state was observed (line 7). The target policies $\boldsymbol{\pi}$ are trained directly using the data within the replay buffer \mathcal{D} . Every N episodes, a new restricted space and goal g is chosen (line 10; see Sec. 3.2 for more). Exploration policies are continuously trained to reach the selected goal (Sec. 3.1).

Algorithm 3: Select Restricted Space and Shared Goal (SelectRestrictedSpaceGoal)

input : counters \mathbf{c} , space tree T_{space} , replay buffer \mathcal{D} , episode

output: selected goal g

- 1 Compute utility of restricted spaces in T_{space}
- 2 Sample a restricted space \mathcal{S}_{k^*} from T_{space} following Eq. (4)
- 3 Sample a batch $B = \{s_i\}_{i=1}^{|B|}$ from \mathcal{D}
- 4 $g = \arg \min_{s \in B} c_{k^*}(\text{proj}_{k^*}(s))$
- 5 **if** $episode \bmod N' = 0$ **then**
- 6 $\text{ExpandSpaceTree}(\mathbf{c}, T_{\text{space}}, k^*) \triangleright \text{Sec. 3.2}$
- 7 **end**
- 8 **return** g

3.1. Training of Exploration Policies

To encourage that exploration policies scout environments in a coordinated manner, we train the exploration policies $\boldsymbol{\mu}$ with an additional modified reward \hat{r} . This modified reward emphasizes the goal g of the exploration. For example, in the two-agent Push-Box task, we use a particular joint location of both agents and the box as a goal. Note, the agents receive a bonus reward \hat{r} when the shared goal g , *i.e.*, the specified state, is reached. The algorithm for training exploration policies is summarized in Alg. 2: standard policy training with a modified reward.

The goal g is obtained via a bottom-up search method. We first explore low-dimensional restricted spaces that are ‘under-explored.’ We discuss this shared goal and restricted space selection method next.

3.2. Shared Goal and Restricted Space Selection

Since the size of the state space \mathcal{S} grows exponentially with the number of agents, conventional exploration strategies (Burda et al., 2019; Mahajan et al., 2019; Ronen

I. Brafman, 2002) which strive to uniformly visit all states are no longer tractable. To address this issue, we propose to first project the state space to restricted spaces, and then perform shared goal driven coordinated exploration in those restricted spaces. For simplicity, we first assume the state space is finite and discrete. We discuss how to extend CMAE to continuous state spaces in Sec. 3.3. In the following, we first show how to select the goal g given a selected restricted space. Then we discuss how to select restricted spaces and expand the space tree T_{space} .

Shared Goal Selection: Given a restricted space \mathcal{S}_{k^*} and its associated counter c_{k^*} , we choose the goal state g by first uniformly sampling a batch of states B from the replay buffer \mathcal{D} . From those states, we select the state with the smallest count as the goal state g , *i.e.*,

$$g = \arg \min_{s \in B} c_{k^*}(\text{proj}_{k^*}(s)), \quad (2)$$

where $\text{proj}_{k^*}(s)$ is the projection from state space to the restricted space \mathcal{S}_{k^*} . To make this concrete, consider again the 2-agent Push-Box game. A restricted space may consist of only the box’s location. Then, from batch B , a state in which the box is in a rarely seen location will be selected as the goal. For each restricted space \mathcal{S}_{k^*} , the associated counter $c_{k^*}(s_{k^*})$ stores the number of times the state s_{k^*} occurred in the low-dimensional restricted space \mathcal{S}_{k^*} .

Given a goal state g , we train exploration policies μ using the method presented in Sec. 3.1 (Alg. 2).

Restricted Space Selection: For an M -dimensional state space, the number of restricted spaces is equivalent to the size of the power set, *i.e.*, 2^M . It’s intractable to study all.

To address this issue, we propose a bottom-up tree-search mechanism to select under-explored restricted spaces. We start from low-dimensional restricted spaces and then gradually grow the search tree to explore higher-dimensional restricted spaces. Specifically, we maintain a space tree T_{space} where each node in the tree represents a restricted space. Each restricted space k in the tree is associated with a utility value u_k , which guides the selection.

The utility permits to identify the under-explored restricted spaces. For this, we study a normalized-entropy-based mechanism to compute the utility of each restricted space in T_{space} . Intuitively, under-explored restricted spaces have lower normalized entropy. To estimate the normalized entropy of a restricted space \mathcal{S}_k , we normalize the counter c_k to obtain a probability distribution $p_k(\cdot) = c_k(\cdot) / \sum_{s \in \mathcal{S}_k} c_k(s)$, which is then used to compute the normalized entropy

$$\eta_k = H_k / H_{\max, k} = - \left(\sum_{s \in \mathcal{S}_k} p_k(s) \log p_k(s) \right) / \log(|\mathcal{S}_k|). \quad (3)$$

Then the utility u_k is given by $u_k = -\eta_k$. Finally, we sample a restricted space k^* following a categorical distribution over all spaces in the space tree T_{space} , *i.e.*,

$$k^* \sim \text{Cat}(\text{softmax}((u_k)_{\mathcal{S}_k \in T_{\text{space}}}).) \quad (4)$$

The restricted space and goal selection method is summarized in Alg. 3. Note that the actual value of $|\mathcal{S}_k|$ is usually unavailable. We defer the details of estimating $|\mathcal{S}_k|$ from observed data to Sec. E.1 in the appendix.

Space Tree Expansion: The space tree T_{space} is initialized with one-dimensional restricted spaces. To consider restricted spaces of higher dimension, we grow the space tree from the current selected restricted space \mathcal{S}_{k^*} every N' episodes. If \mathcal{S}_{k^*} is l -dimensional, we add restricted spaces that are $(l+1)$ -dimensional and contain \mathcal{S}_{k^*} as child nodes of \mathcal{S}_{k^*} . Formally, we initialize the space tree T_{space}^0 via

$$T_{\text{space}}^0 = \{\mathcal{S}_k : |k| = 1, k \in P(\{1, \dots, M\})\}, \quad (5)$$

where M denotes the dimension of the full state space and P denotes the power set. Let $T_{\text{space}}^{(h)}$ and \mathcal{S}_{k^*} denote the space tree after the h -th expansion and the current selected restricted space respectively. Note, \mathcal{S}_{k^*} is sampled according to Eq. (4) and no domain knowledge is used for selecting \mathcal{S}_{k^*} . The space tree after the $(h+1)$ -th expansion is

$$T_{\text{space}}^{(h+1)} = T_{\text{space}}^h \cup \{\mathcal{S}_k : |k| = |k^*| + 1, k^* \subset k, k \in P(\{1, \dots, M\})\}, \quad (6)$$

i.e., all restricted spaces that are $(|k^*|+1)$ -dimensional and contain \mathcal{S}_{k^*} are added. The counters associated with the new restricted spaces are initialized from states in the replay buffer. Specifically, for each newly added restricted space \mathcal{S}_{k^*} , we initialized the corresponding counter to be

$$c_{k^*}(s_{k^*}) = \sum_{s \in \mathcal{D}} \mathbb{1}[\text{proj}_{k^*}(s) = s_{k^*}], \quad (7)$$

where $\mathbb{1}[\cdot]$ is the indicator function (1 if argument is true; 0 otherwise). Once a restricted space was added, we successively increment the counter, *i.e.*, the counter c_{k^*} isn’t recomputed from scratch every episode. This ensures that updating of counters is efficient. Goal selection, space selection and tree expansion are summarized in Alg. 3.

3.3. Counting in Continuous State Spaces

For high-dimensional continuous state spaces, the counters in CMAE could be implemented using counting methods such as neural density models (Ostrovski et al., 2017; Bellemare et al., 2016) or hash-based counting (Tang et al., 2017). Both approaches have been shown to be effective for counting in continuous state spaces.

In our implementation, we adopt hash-based counting (Tang et al., 2017). Hash-based counting discretizes the

state space via a hash function $\phi(s) : S \rightarrow \mathbb{Z}$, which maps a given state s to an integer that is used to index into a table. Specifically, in CMAE, each restricted space \mathcal{S}_k is associated with a hash function $\phi_k(s) : \mathcal{S}_k \rightarrow \mathbb{Z}$, which maps the continuous s_k to an integer. The hash function ϕ_k is used when CMAE updates or queries the counter associated with restricted space \mathcal{S}_k . Empirically, we found CMAE with hash-counting to perform well in environments with continuous state spaces.

3.4. Analysis

To provide more insights into how the proposed method improves data efficiency, we analyze the two major components of CMAE: (1) shared goal exploration and (2) restricted space exploration on a simple multi-player matrix game. We first define the multi-player matrix game:

Example 1. *In a cooperative 2-player l -action matrix game (Myerson, 2013), a payoff matrix $U \in \mathbb{R}^{l \times l}$, which is unobservable to the players, describes the payoff the agents obtain after an action configuration is executed. The agents’ goal is to find the action configuration that maximizes the collective payoff.*

To efficiently find the action configuration that results in maximal payoff, the agents need to uniformly try different action configurations. We show that exploration with shared goals enables agents to see all distinct action configurations more efficiently than exploration without a shared goal. Specifically, when exploring without shared goal, the agents don’t coordinate their behavior. It is equivalent to uniformly picking one action configuration from all configurations. When performing exploration with a shared goal, the least visited action configuration will be chosen as the shared goal. The two agents coordinate to choose the actions that achieve the goal at each step, making exploration more efficient. The following claim formalizes this:

Claim 1. *Consider the 2-player l -action matrix game in Example 1. Let $m = l^2$ denote the total number of action configurations. Let T_m^{share} and $T_m^{\text{non-share}}$ denote the number of steps needed to see all m action configurations at least once for exploration with shared goal and for exploration without shared goal respectively. Then we have $\mathbb{E}[T_m^{\text{share}}] = m$ and $\mathbb{E}[T_m^{\text{non-share}}] = m \sum_{i=1}^m \frac{1}{i} = \Theta(m \ln m)$.¹*

Proof. See supplementary material. \square

Next, we show that whenever the payoff matrix depends only on one agent’s action the expected number of steps to see the maximal reward can be further reduced by first exploring restricted spaces.

Claim 2. *Consider a special case of Example 1 where the payoff matrix depends only on one agent’s action. Let T^{sub} denote the number of steps needed to discover the maximal reward when exploring the action space of agent one and agent two independently. Let T^{full} denote the number of steps needed to discover the maximal reward when the full action space is explored. Then, we have $T^{\text{sub}} = \mathcal{O}(l)$ and $T^{\text{full}} = \mathcal{O}(l^2)$.*

Proof. See supplementary material. \square

Suppose the payoff matrix depends on all agents’ actions. In this case, CMAE will move to the full space after the restricted spaces are well-explored. For this, the expected total number of steps to see the maximal reward is $\mathcal{O}(l + l^2) = \mathcal{O}(l^2)$.

4. Experimental Results

We evaluate CMAE on two challenging environments: (1) a discrete version of the multiple-particle environment (MPE) (Lowe et al., 2017; Wang et al., 2020); and (2) the Starcraft multi-agent challenge (SMAC) (Samvelyan et al., 2019). In both environments, we consider both dense-reward and sparse-reward settings. In a sparse-reward setting, agents don’t receive any intermediate rewards, *i.e.*, agents only receive a reward when a task is completed.

Tasks: We first consider the following tasks of the sparse-reward MPE environment:

Pass-sparse: Two agents operate within two rooms of a 30×30 grid. There is one switch in each room. The rooms are separated by a door and agents start in the same room. The door will open only when one of the switches is occupied. The agents see collective positive reward and the episode terminates only when both agents changed to the other room. The state vector contains x, y locations of all agents and binary variables to indicate if doors are open.

Secret-Room-sparse: Secret-Room extends *Pass*. There are two agents and four rooms. One large room on the left and three small rooms on the right. There is one door between each small room and the large room. The switch in the large room controls all three doors. The switch in each small room only controls the room’s door. The agents need to navigate to one of the three small rooms, *i.e.*, the target room, to receive positive reward. The grid size is 25×25 . The task is considered solved if both agents are in the target room. The state vector contains x, y locations of all agents and binary variables to indicate if doors are open.

Push-Box-sparse: There are two agents and one box in a 15×15 grid. Agents need to push the box to the wall to receive positive reward. The box is heavy, so both agents need to push the box in the same direction at the same time

¹ $\Theta(g)$ means asymptotically bounded above and below by g .

	CMAE (Ours)	Q-learning	Q-learning + Bonus	EITI	EDTI
Pass-sparse	1.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
Secret-Room-sparse	1.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
Push-Box-sparse	1.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
Pass-dense	5.00±0.00	1.25±0.02	1.42±0.14	0.00±0.00	0.18±0.01
Secret-Room-dense	4.00±0.57	1.62±0.16	1.53±0.04	0.00±0.00	0.00±0.00
Push-Box-dense	1.38±0.21	1.58±0.14	1.55±0.04	0.10±0.01	0.05±0.03

Table 1. Final metric of episode rewards of CMAE and baselines on sparse-reward (top) and dense-reward (bottom) MPE tasks.

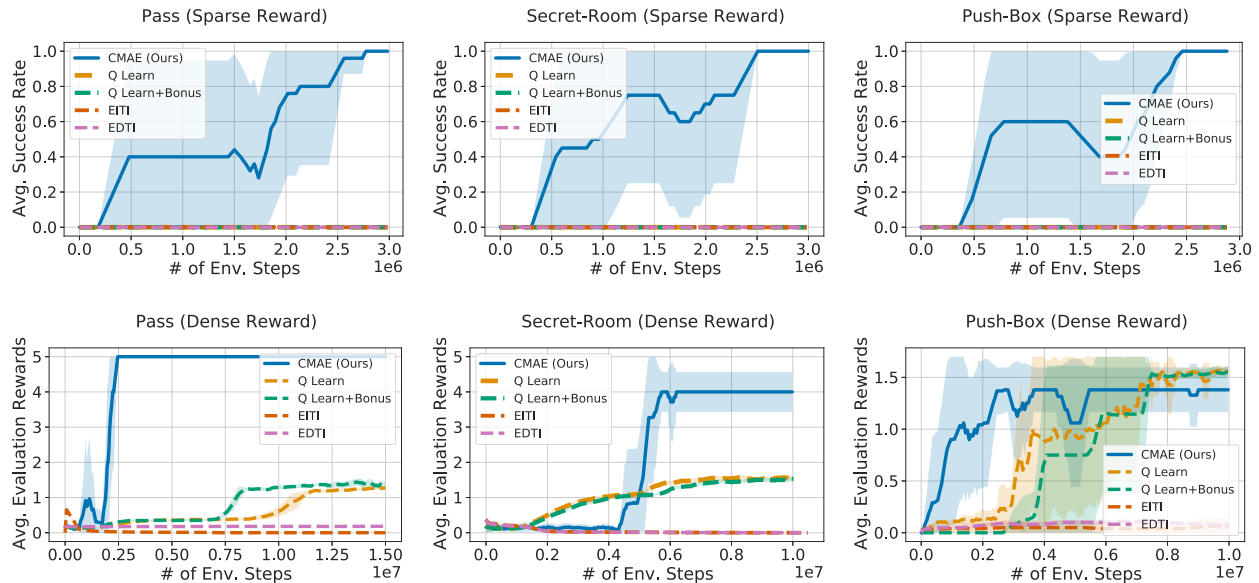


Figure 1. Training curves on sparse-reward and dense-reward MPE tasks.

to move the box. The task is considered solved if the box is pushed to the wall. The state vector contains x, y locations of all agents and the box.

For further details on the sparse-reward MPE tasks, please see Wang et al. (2020). For completeness, in addition to the aforementioned sparse-reward setting, we also consider a dense-reward version of the three tasks. Please see Appendix C for more details on the environment settings.

To evaluate CMAE on environments with continuous state space, we consider three standard tasks in SMAC (Samvelyan et al., 2019): $3m$, $2m_vs_1z$, and $3s_vs_5z$. While the tasks are considered challenging, the commonly used reward is dense, *i.e.*, carefully *hand-crafted* intermediate rewards are used to guide the agents’ learning. However, in many real-world applications, designing effective intermediate rewards may be very difficult or infeasible. Therefore, in addition to the dense-reward setting, we also consider the sparse-reward setting specified by the SMAC environment (Samvelyan et al., 2019) for the three tasks. In SMAC, the state vector contains for all units

on the map: x, y locations, health, shield, and unit type. Note SMAC tasks are partially observable, *i.e.*, agents only observe information of units within a range. Please see Appendix D for more details on the SMAC environment.

Experimental Setup: For MPE tasks, we combine CMAE with Q-learning (Sutton & Barto, 2018; Mnih et al., 2013; 2015). We compare CMAE with exploration via information-theoretic influence (EITI) and exploration via decision-theoretic influence (EDTI) (Wang et al., 2020). EITI and EDTI results are obtained using the publicly available code released by the authors.

For a more complete comparison, we also show the results of Q-learning with ϵ -greedy and Q-learning with count-based exploration (Tang et al., 2017), where exploration bonus is given when a novel state is visited.

For SMAC tasks, we combine CMAE with QMIX (Rashid et al., 2018). We compare with QMIX (Rashid et al., 2018), QMIX with count-based exploration, weighted QMIX (Rashid et al., 2020a), and weighted QMIX with count-based exploration (Tang et al., 2017). For QMIX

	CMAE (Ours)	Weighted QMIX	Weighted QMIX + Bonus	QMIX	QMIX + Bonus
3m-sparse	47.7±35.1	2.7±5.1	11.5±8.6	0.0±0.0	11.7±16.9
2m_vs_1z-sparse	44.3±20.8	0.0±0.0	19.4±18.1	0.0±0.0	19.8±14.1
3s_vs_5z-sparse	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
3m-dense	98.7±1.7	98.3±2.5	98.9±1.7	97.9±3.6	97.3±3.0
2m_vs_1z-dense	98.2±0.1	98.5±0.1	96.0±1.8	97.1±2.4	95.8±1.7
3s_vs_5z-dense	81.3±16.1	92.2±6.6	95.3±2.2	75.0±17.6	78.1±24.4

Table 2. Final metric of success rate (%) of CMAE and baselines on sparse-reward (top) and dense-reward (bottom) SMAC tasks.

and weighted QMIX, we use the publicly available code released by the authors. In all experiments we use restricted spaces of less than four dimensions.

Note, to increase efficiency of the baselines with count-based exploration, in both MPE and SMAC experiments, the counts are shared across all agents. We use ‘+Bonus’ to refer to a baseline with count-based exploration.

Evaluation Protocol: To ensure a rigorous and fair evaluation, we follow the evaluation protocol suggested by Henderson et al. (2017); Colas et al. (2018). We evaluate the target policies in an independent evaluation environment and report *final metric*. The *final metric* is an average episode reward or success rate over the last 100 evaluation episodes, *i.e.*, 10 episodes for each of the last ten policies during training. We repeat all experiments using five runs with different random seeds.

Note that EITI and EDTI (Wang et al., 2020) report the episode rewards *vs.* the number of model updates as an evaluation metric. This isn’t common when evaluating RL algorithms as this plot doesn’t reflect an RL approach’s data efficiency. In contrast, the episode reward *vs.* number of environment steps is a more common metric for data efficiency and is adopted by many RL works (Lillcrap et al., 2016; Wu et al., 2017; Mnih et al., 2013; 2015; Andrychowicz et al., 2017; Shen et al., 2020; Liu et al., 2019a; Li et al., 2019), particularly works on RL exploration (Mahajan et al., 2019; Taiga et al., 2020; Pathak et al., 2017; Tang et al., 2017; Rashid et al., 2020b). Therefore, following most prior works, we report the episode rewards *vs.* the number of environment steps.

Results: We first compare CMAE with baselines on *Pass*, *Secret Room*, and *Push-Box*. The *final metric* and standard deviation is reported in Tab. 1. In the sparse-reward setting, only CMAE is able to solve the tasks, while all baselines do not learn a meaningful policy. Corresponding training curves with standard deviation are included in Fig. 1. CMAE achieves a 100% success rate on *Pass*, *Secret-room*, and *Push-Box* within 3M environment steps. In contrast, baselines cannot solve the task within the given step budget of 3M steps.

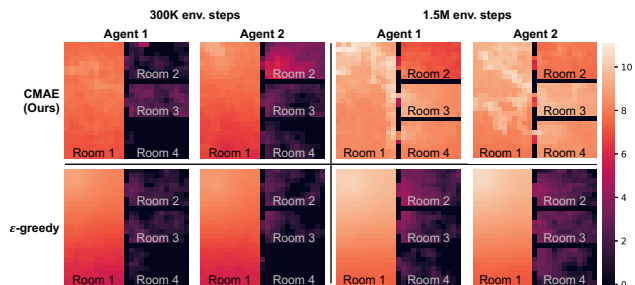


Figure 2. Visitation map (log scale) of CMAE (top) and ϵ -greedy (bottom) on the *Secret-Room* task.

Recently, Taiga et al. (2020) pointed out that many existing exploration strategies excel in challenging sparse-reward tasks but fail in simple tasks that can be solved by using classical exploration methods such as ϵ -greedy. To ensure CMAE doesn’t fail in simpler tasks, we run experiments on the dense-reward version of the three tasks. As shown in Tab. 1, CMAE achieves similar or better performance than the baselines in dense-reward settings.

We also compare the exploration behavior of CMAE to Q-learning with ϵ -greedy exploration using the *Secret-Room* environment. The visit count (in log scale) of each location is visualized in Fig. 2. In early stages of training, both CMAE (top) and ϵ -greedy (bottom) explore only locations in the left room. However, after 1.5M steps, CMAE agents frequently visit the three rooms on the right while ϵ -greedy agents mostly remain within the left room.

On SMAC, we first compare CMAE with baselines in the sparse-reward setting. Since the number of nodes in the space tree grows combinatorially, discovering useful high-dimensional restricted spaces for tasks with high-dimensional state space, such as SMAC, may be infeasible. However, we found empirically that exploring of low-dimensional restricted spaces is already beneficial in a subset of SMAC tasks. The results on SMAC tasks are summarized in Tab. 2, where final metric and standard deviation of evaluation success rate is reported. As shown in Tab. 2 (top), in *3m-sparse* and *2m_vs_1z-sparse*, QMIX and weighted QMIX, which rely on ϵ -greedy exploration, rarely solve the task. When combined with count-based




Figure 3. Training curves on sparse-reward SMAC tasks.

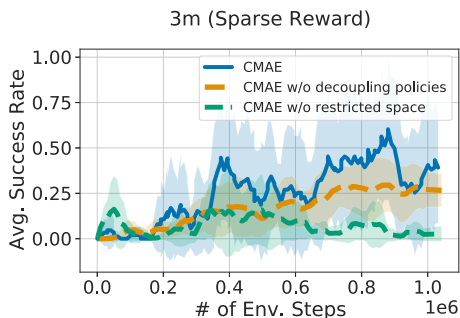


Figure 4. Ablation: CMAE, CMAE without decoupling target and exploration policies, and CMAE without restricted space exploration on *3m-sparse*.

exploration, both QMIX and weighted QMIX are able to achieve 18% to 20% success rate. CMAE achieves much higher success rate of 47.7% and 44.3% on *3m-sparse* and *2m_vs_1z-sparse*, respectively. Corresponding training curves with standard deviation are included in Fig. 3. We also run experiments on dense-reward SMAC tasks, where handcrafted intermediate rewards are available. As shown in Tab. 2 (bottom), CMAE achieves similar performance to state-of-the-art baselines in dense-reward SMAC tasks.

Limitations: To show limitations of the proposed method, we run experiments on the sparse-reward version of *3s_vs_5z*, which is classified as ‘hard’ even in the dense-reward setting (Samvelyan et al., 2019). As shown in Tab. 2 and Fig. 3, CMAE as well as all baselines fail to solve the task. In *3s_vs_5z*, the only winning strategy is to force the enemies to scatter around the map and attend to them one by one (Samvelyan et al., 2019). Without hand-crafted intermediate reward, we found it to be extremely challenging for any approach to pick up this strategy. This demonstrates that efficient exploration for MARL in sparse-reward settings is still a very challenging and open problem, which requires more attention from the community.

Ablation Study: To better understand the approach, we perform an ablation study to examine the effectiveness of

the proposed (1) target and exploration policy decoupling and (2) restricted space exploration. We conduct the experiments on *3m-sparse*. As Fig. 4 shows, without decoupling the exploration and target policies, the success rate drops from 47.7% to 25.4%. In addition, without restricted space exploration, *i.e.*, by directly exploring the full state space, the success rate drops to 9.4%. This demonstrates that the restricted space exploration and policy decoupling are essential to CMAE’s success.

5. Related Work

We discuss recently developed methods for exploration in reinforcement learning, multi-agent reinforcement learning, and concurrent reinforcement learning subsequently.

Exploration for Deep Reinforcement Learning: A wide variety of exploration techniques for deep reinforcement learning have been studied, deviating from classical noise-based methods. Generalization of count-based approaches, which give near-optimal results in tabular reinforcement learning, to environments with continuous state spaces have been proposed. For instance, Bellemare et al. (2016) propose a density model to measure the agent’s uncertainty. Pseudo-counts are derived from the density model which give rise to an exploration bonus encouraging assessment of rarely visited states. Inspired by Bellemare et al. (2016), Ostrovski et al. (2017) discussed a neural density model, to estimate the pseudo count, and Tang et al. (2017) use a hash function to estimate the count.

Besides count-based approaches, meta-policy gradient (Xu et al., 2018) uses the target policy’s improvement as the reward to train an exploration policy. The resulting exploration policy differs from the actor policy, and enables more global exploration. Stadie et al. (2016) propose an exploration strategy based on assigning an exploration bonus from a concurrently learned environment model. Lee et al. (2020) cast exploration as a state marginal matching (SMM) problem and aim to learn a policy for which the state marginal distribution matches a uniform distribution. Other related works on exploration include curiosity-driven

exploration (Pathak et al., 2017), diversity-driven exploration (Hong et al., 2018), GEP-PG (Colas et al., 2018), EX^2 (Fu et al., 2017), bootstrap DQN (Osband et al., 2016) and random network distillation (Burda et al., 2019). In contrast to our approach, all the techniques mentioned above target single-agent deep reinforcement learning.

Multi-agent Deep Reinforcement Learning (MARL): MARL (Lowe et al., 2017; Foerster et al., 2017; Liu et al., 2019b; Rashid et al., 2020a; Jain et al., 2020; Zhou et al., 2020; Christianos et al., 2020; Liu et al., 2020; Jain et al., 2021; Hu et al., 2021) has drawn much attention recently. MADDPG (Lowe et al., 2017) uses a central critic that considers other agents’ action policies to handle the non-stationary environment issues in the multi-agent setting. DIAL (Foerster et al., 2016) uses an end-to-end differentiable architecture that allows agents to learn to communicate. Jiang & Lu (2018) propose an attentional communication model that learns when communication is helpful for a cooperative setting. Foerster et al. (2017) add a ‘fingerprint’ to each transition tuple in the replay memory to track the age of the transition tuple and stabilize training. In ‘Self-Other-Modeling’ (SOM) (Raileanu et al., 2018), an agent uses its own policy to predict other agents’ behavior and states.

While inter-agent communication (Inala et al., 2020; Rangwala & Williams, 2020; Zhang et al., 2020; Ding et al., 2020; Jiang & Lu, 2018; Foerster et al., 2016; Rashid et al., 2018; Omidshafiei et al., 2017; Jain et al., 2019) has been considered, for exploration, multi-agent approaches rely on classical noise-based exploration. As discussed in Sec. 1, a noise-based approach prevents the agents from sharing their understanding of the environment. A team of cooperative agents with a noise-based exploration policy can only explore local regions that are close to their individual actor policy, which contrasts the approach from CMAE.

Recently, approaches that consider coordinated exploration have been proposed. Multi-agent variational exploration (MAVEN) (Mahajan et al., 2019) introduces a latent space for hierarchical control. Agents condition their behavior on the latent variable to perform committed exploration. Influence-based exploration (Wang et al., 2020) captures the influence of one agent’s behavior on others. Agents are encouraged to visit ‘interaction points’ that will change other agents’ behaviour.

Concurrent Deep Reinforcement Learning: Dimakopoulou & Roy (2018) study coordinated exploration in concurrent reinforcement learning, maintaining an environment model and extending posterior sampling such that agents explore in a coordinated fashion. Parisotto et al. (2019) proposed concurrent meta reinforcement learning (CMRL) which permits a set of parallel agents to communicate with each other and find efficient exploration

strategies. The concurrent setting differs from the multi-agent setting of our approach. In a concurrent setting, agents operate in different instances of an environment, *i.e.*, one agent’s action has no effect on the observations and rewards received by other agents. In contrast, in the multi-agent setting, agents share the same instance of an environment. An agent’s action changes observations and rewards observed by other agents.

6. Conclusion

We propose cooperative multi-agent exploration (CMAE). It defines shared goals and learns coordinated exploration policies. To find a goal for efficient exploration we study restricted space selection which helps, particularly in sparse-reward environments. Empirically, we demonstrate that CMAE increases exploration efficiency. We hope this is a first step toward efficient coordinated exploration.

Acknowledgement: This work is supported in part by NSF under Grant #1718221, 2008387, 2045586, and MRI #1725729, UIUC, Samsung, Amazon, 3M, and Cisco Systems Inc. RY is supported by a Google Fellowship.

References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., and Zaremba, W. Hindsight experience replay. In *Proc. NeurIPS*, 2017.
- Bazzan, A. L. C. Opportunities for multiagent systems and multiagent reinforcement learning in traffic control. In *Proc. AAMAS*, 2008.
- Bellemare, M. G., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Proc. NeurIPS*, 2016.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. In *Proc. ICLR*, 2019.
- Christianos, F., Schafer, L., and Albrecht, S. V. Shared experience actor-critic for multi-agent reinforcement learning. In *Proc. NeurIPS*, 2020.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *arXiv*, 2014.
- Colas, C., Sigaud, O., and Oudeyer, P.-Y. GEP-PG: Decoupling exploration and exploitation in deep reinforcement learning algorithms. In *Proc. ICML*, 2018.

- Dimakopoulou, M. and Roy, B. V. Coordinated exploration in concurrent reinforcement learning. In *Proc. ICML*, 2018.
- Ding, Z., Huang, T., and Lu, Z. Learning individually inferred communication for multi-agent cooperation. In *Proc. NeurIPS*, 2020.
- Foerster, J., Nardelli, N., Farquhar, G., Afouras, T., Torr, P. H. S., Kohli, P., and Whiteson, S. Stabilising experience replay for deep multi-agent reinforcement learning. In *Proc. ICML*, 2017.
- Foerster, J., Farquhar, G., Afouras, T., Nardelli, N., and Whiteson, S. Counterfactual multi-agent policy gradients. In *Proc. AAAI*, 2018.
- Foerster, J. N., Assael, Y. M., de Freitas, N., and Whiteson, S. Learning to communicate with deep multi-agent reinforcement learning. In *Proc. ICML*, 2016.
- Fu, J., Co-Reyes, J. D., and Levine, S. Ex2: Exploration with exemplar models for deep reinforcement learning. In *Proc. NeurIPS*, 2017.
- Hausknecht, M. and Stone, P. Deep recurrent q-learning for partially observable mdps. In *arXiv*, 2015.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. Deep reinforcement learning that matters. In *Proc. AAAI*, 2017.
- Hong, Z.-W., Shann, T.-Y., Su, S.-Y., Chang, Y.-H., and Lee, C.-Y. Diversity-driven exploration strategy for deep reinforcement learning. In *Proc. NeurIPS*, 2018.
- Hu, S., Zhu, F., Chang, X., and Liang, X. Updet: Universal multi-agent rl via policy decoupling with transformers. In *Proc. ICLR*, 2021.
- Hüttenrauch, M., Šošić, A., and Neumann, G. Deep reinforcement learning for swarm systems. *JMLR*, 2019.
- Inala, J. P., Yang, Y., Paulos, J., Pu, Y., Bastani, O., Kumar, V., Rinard, M., and Solar-Lezama, A. Neurosymbolic transformers for multi-agent communication. In *Proc. NeurIPS*, 2020.
- Jain, U., Weihs, L., Kolve, E., Rastegari, M., Lazebnik, S., Farhadi, A., Schwing, A., and Kembhavi, A. Two body problem: Collaborative visual task completion. In *Proc. CVPR*, 2019.
- Jain, U., Weihs, L., Kolve, E., Farhadi, A., Lazebnik, S., Kembhavi, A., and Schwing, A. G. A cordial sync: Going beyond marginal policies for multi-agent embodied tasks. In *Proc. ECCV*, 2020.
- Jain, U., Liu, I.-J., Lazebnik, S., Kembhavi, A., Weihs, L., and Schwing, A. Gridtopix: Training embodied agents with minimal supervision. In *arXiv*, 2021.
- Jiang, J. and Lu, Z. Learning attentional communication for multi-agent cooperation. In *Proc. NeurIPS*, 2018.
- Lee, L., Eysenbach, B., Parisotto, E., Xing, E., Levine, S., and Salakhutdinov, R. Efficient exploration via state marginal matching. In *arXiv*, 2020.
- Li, Y., Liu, I.-J., Yuan, Y., Chen, D., Schwing, A., and Huang, J. Accelerating distributed reinforcement learning with in-switch computing. In *Proc. ISCA*, 2019.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *Proc. ICLR*, 2016.
- Liu, I.-J., Peng, J., and Schwing, A. G. Knowledge flow: Improve upon your teachers. In *Proc. ICLR*, 2019a.
- Liu, I.-J., Yeh, R. A., and Schwing, A. G. PIC: permutation invariant critic for multi-agent deep reinforcement learning. In *Proc. CoRL*, 2019b.
- Liu, I.-J., Yeh, R. A., and Schwing, A. G. High-throughput synchronous deep rl. In *Proc. NeurIPS*, 2020.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proc. NeurIPS*, 2017.
- Mahajan, A., Rashid, T., Samvelyan, M., and Whiteson, S. MAVEN: multi-agent variational exploration. In *Proc. NeurIPS*, 2019.
- Matignon, L., j. Laurent, G., and fort piat, N. L. Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. *The Knowledge Engineering Review*, 2012.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. In *NeurIPS Deep Learning Workshop*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 2015.
- Myerson, R. B. *Game Theory*. Harvard University Press, 2013.

- Omidshafiei, S., Pazis, J., Amato, C., How, J. P., and Vian, J. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *Proc. ICML*, 2017.
- Osband, I., Blundell, C., Pritzel, A., and Roy, B. V. Deep exploration via bootstrapped dqn. In *Proc. NeurIPS*, 2016.
- Ostrovski, G., Bellemare, M. G., van den Oord, A., and Munos, R. Count-based exploration with neural density models. In *Proc. ICML*, 2017.
- Parisotto, E., Ghosh, S., Yalamanchi, S. B., Chinnabireddy, V., Wu, Y., and Salakhutdinov, R. Concurrent meta reinforcement learning. In *arXiv*, 2019.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proc. ICML*, 2017.
- Raileanu, R., Denton, E., Szlam, A., and Fergus, R. Modeling others using oneself in multi-agent reinforcement learning. In *Proc. ICML*, 2018.
- Rangwala, M. and Williams, R. Learning multi-agent communication through structured attentive reasoning. In *Proc. NeurIPS*, 2020.
- Rashid, T., Samvelyan, M., de Witt, C. S., Farquhar, G., Foerster, J., and Whiteson, S. QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proc. ICML*, 2018.
- Rashid, T., Farquhar, G., Peng, B., and Whiteson, S. Weighted qmix: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proc. NeurIPS*, 2020a.
- Rashid, T., Peng, B., Böhrer, W., and Whiteson, S. Optimistic exploration even with a pessimistic initialisation. In *Proc. ICLR*, 2020b.
- Ronen I. Brafman, M. T. R-max - a general polynomial time algorithm for near-optimal reinforcement learning. In *JMLR*, 2002.
- Samvelyan, M., Rashid, T., de Witt, C. S., Farquhar, G., Nardelli, N., Rudner, T. G. J., Hung, C.-M., Torr, P. H. S., Foerster, J., and Whiteson, S. The starcraft multi-agent challenge. In *arxiv*, 2019.
- Shen, J., Zhao, H., Zhang, W., and Yu, Y. Model-based policy optimization with unsupervised model adaptation. In *Proc. NeurIPS*, 2020.
- Stadie, B. C., Levine, S., and Abbeel, P. Incentivizing exploration in reinforcement learning with deep predictive models. In *Proc. ICLR*, 2016.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., and Graepel, T. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proc. AAMAS*, 2018.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. The MIT Press, 2018.
- Swamy, G., Reddy, S., Levine, S., and Dragan, A. D. Scaled autonomy: Enabling human operators to control robot fleets. In *Proc. ICRA*, 2020.
- Taiga, A. A., Fedus, W., Machado, M. C., Courville, A., and Bellemare, M. G. On bonus based exploration methods in the arcade learning environment. In *Proc. ICLR*, 2020.
- Tampuu, A., Matiisen, T., Kodelja, D., Kuzovkin, I., Korjus, K., Aru, J., Aru, J., and Vicente, R. Multiagent cooperation and competition with deep reinforcement learning. In *arxiv*, 2015.
- Tan, M. Multiagent reinforcement learning independent vs cooperative agents. In *Proc. ICML*, 1993.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, X., Duan, Y., Schulman, J., Turck, F. D., and Abbeel, P. Exploration: A study of count-based exploration for deep reinforcement learning. In *Proc. NeurIPS*, 2017.
- Wang, T., Wang, J., Wu, Y., and Zhang, C. Influence-based multi-agent exploration. In *Proc. ICLR*, 2020.
- Wu, Y., Mansimov, E., Liao, S., Grosse, R., and Ba, J. Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation. In *Proc. NeurIPS*, 2017.
- Xu, T., Liu, Q., Zhao, L., and Peng, J. Learning to explore via meta-policy gradient. In *Proc. ICML*, 2018.
- Yang, Y., Luo, R., Li, M., Zhou, M., Zhang, W., and Wang, J. Mean field multi-agent reinforcement learning. In *Proc. ICML*, 2018.
- Zhang, S. Q., Zhang, Q., and Lin, J. Succinct and robust multi-agent communication with temporal message control. In *Proc. NeurIPS*, 2020.
- Zhou, M., Liu, Z., Sui, P., Li, Y., and Chung, Y. Y. Learning implicit credit assignment for cooperative multi-agent reinforcement learning. In *Proc. NeurIPS*, 2020.