

Offline Decentralized Multi-Agent Reinforcement Learning

Jiechuan Jiang^a and Zongqing Lu^{a,*}

^aPeking University

Abstract. In many real-world multi-agent cooperative tasks, due to high cost and risk, agents cannot continuously interact with the environment and collect experiences during learning, but have to learn from offline datasets. However, the transition dynamics in the dataset of each agent can be much different from the ones induced by the learned policies of other agents in execution, creating large errors in value estimates. Consequently, agents learn uncoordinated low-performing policies. In this paper, we propose a framework for offline decentralized multi-agent reinforcement learning, which exploits *value deviation* and *transition normalization* to deliberately modify the transition probabilities. Value deviation optimistically increases the transition probabilities of high-value next states, and transition normalization normalizes the transition probabilities of next states. They together enable agents to learn high-performing and coordinated policies. Theoretically, we prove the convergence of Q-learning under the altered *non-stationary* transition dynamics. Empirically, we show that the framework can be easily built on many existing offline reinforcement learning algorithms and achieve substantial improvement in a variety of multi-agent tasks.

1 Introduction

Reinforcement learning (RL) has shown great potential in domains, including recommendation systems [12], games [27], and robotics [6]. When applying RL to real-world applications, the agent needs to continuously interact with the environment and collect the experiences for learning, which however can be costly, risky, and time-consuming. One way to address this is offline RL, where the agent learns the policy, without interacting with the environment, from a fixed dataset of experiences collected by any behavior policy. The main challenge of offline RL is the *extrapolation error*, an error incurred by the mismatch between the experience distributions of the learned policy and dataset [5]. Many offline RL methods [5, 9, 30, 10] have been proposed to address the value overestimate caused by the mismatch between action distributions of the learned policy and behavior policy. However, much less attention has been paid to the mismatch between the transition dynamics in the dataset and the real ones in execution. The main reason is that the mismatch of transition dynamics is negligible given a large dataset, if the environment is stationary. Nevertheless, in many real-world scenarios, there are *other* learning agents in the environment, *e.g.*, autonomous driving. That means the policies of other agents during data collection can be significantly different from their policies in execution, which creates the

mismatch of transition dynamics and hence undermines the learned policy.

More specifically, these scenarios can be formulated as the offline decentralized multi-agent setting, where agents cooperate on the task but each learns the policy from its own fixed dataset. The dataset of each agent is independently collected by any behavior policies *arbitrarily* (*i.e.*, we do not make any assumptions on data collection) and contains its own action instead of the joint action of all agents. This setting resembles many industrial applications where agents belong to different companies and the actions of other agents are not accessible, *e.g.*, autonomous vehicles or robots. Apparently, this setting does not fit the paradigm of offline centralized training and decentralized execution (CTDE) in multi-agent reinforcement learning (MARL)¹ [32], and each agent has to learn a cooperative policy in an *offline and fully decentralized* way. In such decentralized multi-agent settings, from the perspective of an individual agent, other agents are a part of the environment, thus the transition dynamics experienced by each agent depend on the policies of other agents and change as other agents update their policies [2]. As the learned policies of other agents will be inconsistent with their behavior policies, the transition dynamics induced by the learned policies of other agents in execution will be different from the transition dynamics in the dataset, causing a large mismatch between the transition dynamics. This mismatch can lead to a low-performing policy for each agent. Furthermore, as agents learn in a decentralized way on the datasets with different distributions of experiences collected by various behavior policies, the estimated values of the same state can be much different between agents, which causes the learned policies cannot well coordinate with each other.

In this paper, we propose a framework for offline fully decentralized multi-agent reinforcement learning, where, to overcome the mismatch between transition dynamics and miscoordination, *value deviation* and *transition normalization* are introduced to deliberately modify the transition probabilities in the dataset. During data collection, if one agent takes a ‘good’ action while other agents take ‘bad’ actions at a state, the transition probabilities of low-value next states will be high. Thus, the Q-value of the ‘good’ action will be *underestimated* by the agent. As other agents are also learning, their learned policies are expected to become better than their behavior policies. Therefore, for each agent, the transition probabilities of high-value next states in execution will be higher than those in the dataset. So, we let each agent be *optimistic* towards other agents and multiply

¹ Even if the datasets of all agents can be accessed in a centralized way, a full dataset that contains the joint actions cannot be constructed, because the datasets are fully independent without any common information.

* Corresponding Author. Email: zongqing.lu@pku.edu.cn

the transition probabilities by the deviation of the value of next state from the expected value over all next states, to make the estimated transition probabilities *close* to the transition probabilities induced by the learned policies of other agents. To address the miscoordination caused by diverse value estimates of agents, we normalize the transition probabilities in the dataset to be uniform. Transition normalization helps to build the consensus about value estimate among agents and hence promotes coordination. By combining value deviation and transition normalization, agents are able to learn high-performing and coordinated policies in an offline and fully decentralized way.

Value deviation and transition normalization make the transition dynamics non-stationary during learning. However, we mathematically prove the convergence of Q-learning under such purposely controlled non-stationary transition dynamics. Moreover, by derivation, we show that value deviation and transition normalization can take effect *merely* as the weights of the objective function, thus our framework can be easily built on many existing offline single-agent RL algorithms that address the overestimation incurred by out-of-distribution actions. Empirically, we provide an example instantiation and the thorough analysis of our framework on BCQ [5], termed MABCQ, and also test the variants on CQL [10] and TD3+BC [4], termed MACQL and MATD3+BC respectively. Experimental results show that our method substantially outperforms baselines in a variety of multi-agent tasks, including multi-agent mujoco [1], SMAC [22] and MPE [14]. To the best of our knowledge, this is the *first* work for offline and decentralized multi-agent reinforcement learning.

2 Preliminaries

We consider N agents in multi-agent MDP [16] $M_{\text{env}} = \langle \mathcal{S}, \mathcal{A}, R, P_{\text{env}}, \gamma \rangle$ with the state space \mathcal{S} and the joint action space \mathcal{A} . At each timestep, each agent i gets state² s , and performs an individual action a_i , and the environment transitions to the next state s' by taking the joint action \mathbf{a} with the transition probability $P_{\text{env}}(s'|s, \mathbf{a})$. All agents get a shared reward $r = R(s)$, which is simplified to just depending on state [23]³. All the agents learn to maximize the expected return $\mathbb{E} \sum_{t=0}^{\infty} \gamma^t r_t$, where γ is a discount factor. However, in the fully decentralized learning, M_{env} is partially observable to each agent since each agent only observes its own action a_i instead of the joint action \mathbf{a} . During execution, from the perspective of each agent i , there is an experienced MDP $M_{\mathcal{E}_i} = \langle \mathcal{S}, \mathcal{A}_i, R, P_{\mathcal{E}_i}, \gamma \rangle$ with the individual action space \mathcal{A}_i and the **online transition probability**

$$P_{\mathcal{E}_i}(s'|s, a_i) = \sum_{\mathbf{a}_{-i}} P_{\text{env}}(s'|s, a_i, \mathbf{a}_{-i}) \pi_{\mathcal{E}_{-i}}(\mathbf{a}_{-i}|s),$$

where \mathbf{a}_{-i} and $\pi_{\mathcal{E}_{-i}}$ respectively denote the joint action and the learned joint policy of all agents except agent i . However, if the agent cannot interact with other agents in the environment, $P_{\mathcal{E}_i}$ is *unknown*.

In *offline decentralized settings*, each agent i only has access to a fixed offline dataset \mathcal{B}_i , which is pre-collected by behavior policies $\pi_{\mathcal{B}}$ and contains the tuples (s, a_i, r, s') . As defined in BCQ [5], the visible MDP $M_{\mathcal{B}_i} = \langle \mathcal{S}, \mathcal{A}_i, R, P_{\mathcal{B}_i}, \gamma \rangle$ is constructed on \mathcal{B}_i , which has the **offline transition probability**

$$P_{\mathcal{B}_i}(s'|s, a_i) = \sum_{\mathbf{a}_{-i}} P_{\text{env}}(s'|s, a_i, \mathbf{a}_{-i}) \pi_{\mathcal{B}_{-i}}(\mathbf{a}_{-i}|s),$$

As the learned policies of other agents $\pi_{\mathcal{E}_{-i}}$ may greatly deviate from their behavior policies $\pi_{\mathcal{B}_{-i}}$, $P_{\mathcal{B}_i}$ can be largely biased from $P_{\mathcal{E}_i}$. We define the discrepancy between $P_{\mathcal{B}_i}$ and $P_{\mathcal{E}_i}$ as **transition bias**. The large extrapolation errors caused by transition bias, and the differences in value estimates between agents lead to uncoordinated low-performing policies.

Table 1: The matrix game.

		Agent 2	
		a_1 (0.4)	a_2 (0.6)
Agent 1	a_1 (0.8)	1	5
	a_2 (0.2)	6	1

Table 2: Transition probabilities and expected returns calculated in the dataset.

Agent 1	action	transition	expected return
	a_1		$p(1 a_1) = 0.4$
		$p(5 a_1) = 0.6$	
a_2		$p(6 a_2) = 0.4$	3.0
		$p(1 a_2) = 0.6$	

Agent 2	action	transition	expected return
	a_1		$p(1 a_1) = 0.8$
		$p(6 a_1) = 0.2$	
a_2		$p(5 a_2) = 0.8$	4.2
		$p(1 a_2) = 0.2$	

To intuitively illustrate the issue, we introduce a two-player matrix game as depicted in Table 1, where the action distributions of the behavior policies of the two agents are $[0.8, 0.2]$ and $[0.4, 0.6]$, respectively. Table 2 shows the transition probabilities and expected returns calculated independently by the agents from the datasets collected by the behavior policies. However, as the behavior policies are poor, when agent 1 chooses the optimal action a_2 , agent 2 has a higher probability to select the suboptimal action a_2 , which leads to low transition probabilities of *true* high-value next states ($p(6|a_2) = 0.4$ vs. $p(1|a_2) = 0.6$). So the agents underestimate the optimal actions and converge to the suboptimal policies (a_1, a_2) , rather than the optimal policy (a_2, a_1) .

3 Offline Decentralized MARL Framework

In fully decentralized learning, as the policies of other agents are not accessible, it is hard for an agent to learn a policy that can coordinate well with other agents, merely from the dataset that only contains its own action. To tackle this challenging problem, we propose a framework, which newly introduces value deviation and transition normalization to address transition bias and miscoordination, and leverages the offline single-agent RL algorithm to avoid out-of-distribution actions. The convergence under the non-stationary transition dynamics is theoretically guaranteed, and an example instantiation of the framework is provided.

3.1 Value Deviation

If the behavior policies of some agents are low-performing during data collection, they usually take ‘bad’ actions to cooperate with the ‘good’ actions of other agents, which leads to high transition probabilities of low-value next states. When agent i performs Q-learning with the dataset \mathcal{B}_i , the Bellman operator \mathcal{T} is approximated by the

² Global state is only for the convenience of theoretical analysis [8] the agents could learn based on partial observation in practice.

³ This simplification is only for the convenience of theoretical analysis, but not necessary in practice.

transition probability $P_{\mathcal{B}_i}(s'|s, a_i)$ to estimate the expectation over s' :

$$\mathcal{T}Q_i(s, a_i) = \mathbb{E}_{s' \sim P_{\mathcal{B}_i}(s'|s, a_i)} \left[r + \gamma \max_{\hat{a}_i} Q_i(s', \hat{a}_i) \right].$$

If $P_{\mathcal{B}_i}$ of a high-value s' is lower than $P_{\mathcal{E}_i}$, the Q-value of the state-action pair (s, a_i) is underestimated, which will cause large extrapolation error.

However, as the policies of other agents are also updating towards maximizing the Q-values, $P_{\mathcal{E}_i}$ of high-value next states will grow higher than $P_{\mathcal{B}_i}$. Thus, we let each agent be optimistic towards other agents and modify $P_{\mathcal{B}_i}$ as

$$P_{\mathcal{B}_i}(s'|s, a_i) \cdot \underbrace{\left(1 + \frac{V_i^*(s') - \mathbb{E}_{s'} V_i^*(s')}{|\mathbb{E}_{s'} V_i^*(s')|}\right)}_{\text{value deviation}} \cdot \frac{1}{z_i^{vd}},$$

where the state value $V_i^*(s) = \max_{a_i} Q_i(s, a_i)$, $1 + \frac{V_i^*(s') - \mathbb{E}_{s'} V_i^*(s')}{|\mathbb{E}_{s'} V_i^*(s')|}$ is the deviation of the value of next state from the expected value over all next states, which increases the transition probabilities of high-value next states and decreases those of low-value next states, and $z_i^{vd} = \sum_{s'} P_{\mathcal{B}_i}(s'|s, a_i) * (1 + \frac{V_i^*(s') - \mathbb{E}_{s'} V_i^*(s')}{|\mathbb{E}_{s'} V_i^*(s')|})$ is a normalization term to make sure the sum of the transition probabilities is one. *Value deviation* modifies the transition probabilities to be close to $P_{\mathcal{E}_i}$ and hence reduces the transition bias. The optimism towards other agents helps the agent discover potential good actions which are hidden by the poor behavior policies of other agents.

3.2 Transition Normalization

As \mathcal{B}_i of each agent is individually collected by different behavior policies, the diverse combinations of behavior policies of all agents lead to the value of the same state s being overestimated by some agents, while being underestimated by others. Since the agents are trained to reach high-value states, the large disagreement on state values will cause miscoordination of the learned policies. To overcome the problem, we normalize $P_{\mathcal{B}_i}$ to be uniform over next states,

$$P_{\mathcal{B}_i}(s'|s, a_i) \cdot \underbrace{\frac{1}{P_{\mathcal{B}_i}(s'|s, a_i)}}_{\text{transition normalization}} \cdot \frac{1}{z_i^{tn}}{}^4$$

where z_i^{tn} is a normalization term that is the number of different s' given (s, a_i) in \mathcal{B}_i . *Transition normalization* enforces that each agent has the same $P_{\mathcal{B}_i}$ when it acts the learned action a_i^* on the same state s , and we have the following proposition.

Proposition 1. *In episodic environments, if each agent i performs Q-learning on \mathcal{B}_i , all agents will converge to the same V^* if they have the same transition probability on any state where each agent i acts the learned action a_i^* .*

The proof is provided in Appendix⁵. This proposition implies that transition normalization can enable agents to have the same state value estimate. However, to satisfy $P_{\mathcal{B}_1}(s'|s, a_1^*) = P_{\mathcal{B}_2}(s'|s, a_2^*) = \dots = P_{\mathcal{B}_N}(s'|s, a_N^*)$ for all $s' \in \mathcal{S}$ in the datasets, the agents should have the same set of s' at (s, a^*) , which is a strong assumption. In practice, although the assumption is not strictly satisfied, transition normalization can still normalize the transition probabilities, encouraging the estimated state value V^* to be close to each other.

⁴ Although the modification can be simply written as $1/z_i^{tn}$, we present it in such a way that we can later conveniently combine transition normalization with value deviation.

⁵ Appendix is available at <https://arxiv.org/abs/2108.01832>.

3.3 Optimization Objective

Combining value deviation $1 + (V_i^*(s') - \mathbb{E}_{s'} V_i^*(s')) / |\mathbb{E}_{s'} V_i^*(s')|$, denoted as λ_{vd_i} , and transition normalization $1/P_{\mathcal{B}_i}(s'|s, a_i)$, denoted as λ_{tn_i} , we modify $P_{\mathcal{B}_i}$ as

$$\hat{P}_{\mathcal{B}_i}(s'|s, a_i) = P_{\mathcal{B}_i}(s'|s, a_i) \cdot \frac{\lambda_{tn_i} \lambda_{vd_i}}{z_i},$$

where $z_i = \sum_{s'} (1 + (V_i^*(s') - \mathbb{E}_{s'} V_i^*(s')) / |\mathbb{E}_{s'} V_i^*(s')|)$ is the normalization term. Indeed, $\hat{P}_{\mathcal{B}_i}$ makes offline learning on \mathcal{B}_i similar to *online* decentralized MARL. In the initial stage, λ_{vd_i} is close to 1 since $Q_i(s, a_i)$ is not learned yet, and the transition probabilities are uniform, resembling the start stage in online learning where all agents are acting randomly. During training, the transition probabilities of high-value states gradually grow by value deviation, which is an analogy of other agents improving their policies in online learning. Therefore, $\hat{P}_{\mathcal{B}_i}$ encourages the agents to learn high-performing policies and improves coordination.

Although $\hat{P}_{\mathcal{B}_i}$ is non-stationary (*i.e.*, λ_{vd_i} changes along with the updates of Q-value), we have the following theorem that guarantees the convergence of Bellman operator \mathcal{T} under $\hat{P}_{\mathcal{B}_i}$,

$$\mathcal{T}Q_i(s, a_i) = \mathbb{E}_{s' \sim \hat{P}_{\mathcal{B}_i}(s'|s, a_i)} \left[r + \gamma \max_{\hat{a}_i} Q_i(s', \hat{a}_i) \right].$$

Theorem 1. *Under the non-stationary transition probability $\hat{P}_{\mathcal{B}_i}$, the Bellman operator \mathcal{T} is a contraction and converges to a unique fixed point when $\gamma < r_{\min} / (2r_{\max} - r_{\min})$, if the reward is bounded by the positive region $[r_{\min}, r_{\max}]$.*

The proof is provided in Appendix. As any positive affine transformation of the reward function does not change the optimal policy in the fixed-horizon environments [35], Theorem 1 holds in general, and we can rescale the reward to make r_{\min} arbitrarily close to r_{\max} so as to make the upper bound of γ close to 1.

In deep reinforcement learning, directly modifying the transition probability is infeasible. However, we can instead modify the sampling probability to achieve the same effect. The optimization objective of decentralized deep Q-learning $\mathbb{E}_{p_{\mathcal{B}_i}(s, a_i, s')} |Q_i(s, a_i) - y_i|^2$ is calculated by sampling the batch from \mathcal{B}_i according to the sampling probability $p_{\mathcal{B}_i}(s, a_i, s')$. By factorizing $p_{\mathcal{B}_i}(s, a_i, s')$, we have

$$\underbrace{p_{\mathcal{B}_i}(s, a_i, s')}_{\text{sampling probability}} = p_{\mathcal{B}_i}(s, a_i) \cdot \underbrace{P_{\mathcal{B}_i}(s'|s, a_i)}_{\text{transition probability}}.$$

Therefore, we can modify the transition probability as $\frac{\lambda_{tn_i} \lambda_{vd_i}}{z_i} P_{\mathcal{B}_i}(s'|s, a_i)$ and scale $p_{\mathcal{B}_i}(s, a_i)$ with z_i . Then, the sampling probability can be re-written as

$$\underbrace{\lambda_{tn_i} \lambda_{vd_i} p_{\mathcal{B}_i}(s, a_i, s')}_{\text{modified sampling probability}} = z_i p_{\mathcal{B}_i}(s, a_i) \cdot \underbrace{\frac{\lambda_{tn_i} \lambda_{vd_i}}{z_i} P_{\mathcal{B}_i}(s'|s, a_i)}_{\text{modified transition probability}}.$$

Since z_i is independent of s' , it can be regarded as a scale factor on $p_{\mathcal{B}_i}(s, a_i)$, which will not change the expected target value y_i . Thus, sampling batches according to the modified sampling probability can achieve the same effect as modifying the transition probability. Using importance sampling, the modified optimization objective is

$$\begin{aligned} & \mathbb{E}_{\lambda_{tn_i} \lambda_{vd_i} p_{\mathcal{B}_i}(s, a_i, s')} |Q_i(s, a_i) - y_i|^2 \\ &= \mathbb{E}_{p_{\mathcal{B}_i}(s, a_i, s')} \frac{\lambda_{tn_i} \lambda_{vd_i} p_{\mathcal{B}_i}(s, a_i, s')}{p_{\mathcal{B}_i}(s, a_i, s')} |Q_i(s, a_i) - y_i|^2 \\ &= \mathbb{E}_{p_{\mathcal{B}_i}(s, a_i, s')} \lambda_{tn_i} \lambda_{vd_i} |Q_i(s, a_i) - y_i|^2. \end{aligned}$$

Algorithm 1 MABCQ

-
- 1: **for** $i \in N$ **do**
 - 2: Initialize the conditional VAEs:
 $G_i^1 = \{E_i^1(\mu^1, \sigma^1 | s, a), D_i^1(a | s, z^1)\},$
 $G_i^2 = \{E_i^2(\mu^2, \sigma^2 | s, a, s'), D_i^2(a | s, s', z^2)\}.$
 - 3: Initialize Q-network Q_i , perturbation network ξ_i , and their target networks \hat{Q}_i and $\hat{\xi}_i$.
 - 4: Fit the VAEs G_i^1 and G_i^2 using \mathcal{B}_i .
 - 5: **for** $t = 1, \dots, \text{max_update}$ **do**
 - 6: Sample a mini-batch from \mathcal{B}_i .
 - 7: Update Q_i by minimizing (1).
 - 8: Update ξ_i by maximizing (2).
 - 9: Update the target networks \hat{Q}_i and $\hat{\xi}_i$.
 - 10: **end for**
 - 11: **end for**
-

We can see that λ_{tn_i} and λ_{vd_i} simply take effect as the weights of the objective function, which makes them easily integrated with existing offline RL methods.

3.4 An Example Instantiation

Our framework can be practically instantiated on many offline single-agent RL algorithms that address the overestimation incurred by out-of-distribution actions. Here, we give the instantiation of the framework on BCQ [5], termed MABCQ. To make MABCQ adapt to high-dimensional continuous spaces, for each agent i , we train a Q-network Q_i , a perturbation network ξ_i , and a conditional VAE $G_i^1 = \{E_i^1(\mu^1, \sigma^1 | s, a), D_i^1(a | s, z^1 \sim (\mu^1, \sigma^1))\}$. In execution, each agent i generates n actions by G_i^1 , adds small perturbations $\in [-\Phi, \Phi]$ on the actions using ξ_i , and then selects the action with the highest value in Q_i . The policy can be written as

$$\pi_i(s) = \underset{a_i^j + \xi_i(s, a_i^j)}{\operatorname{argmax}} Q_i(s, a_i^j + \xi(s, a_i^j)),$$

$$\text{where } \left\{ a_i^j \sim G_i^1(s) \right\}_{j=1}^n.$$

Q_i is updated by minimizing

$$\mathbb{E}_{p_{\mathcal{B}_i}(s, a_i, s')} \lambda_{tn_i} \lambda_{vd_i} |Q_i(s, a_i) - y_i|^2,$$

$$\text{where } y_i = r + \gamma \hat{Q}_i(s', \hat{\pi}_i(s')).$$
(1)

y_i is calculated by the target networks \hat{Q}_i and $\hat{\xi}_i$, where $\hat{\pi}_i$ is correspondingly the policy induced by \hat{Q}_i and $\hat{\xi}_i$. ξ_i is updated by maximizing

$$\mathbb{E}_{p_{\mathcal{B}_i}(s, a_i, s')} \lambda_{tn_i} \lambda_{vd_i} Q_i(s, a_i + \xi_i(s, a_i)).$$
(2)

To estimate λ_{vd_i} , we need $V_i^*(s') = \hat{Q}_i(s', \hat{\pi}_i(s'))$ and $\mathbb{E}_{s'}[V_i^*(s')] = \frac{1}{\gamma}(\hat{Q}_i(s, a_i) - r)$, which can be estimated from the sample without actually going through all s' . We estimate λ_{vd_i} using the target networks to stabilize λ_{vd_i} along with the updates of Q_i and ξ_i . To avoid extreme values, we clip λ_{vd_i} to the region $[1 - \epsilon, 1 + \epsilon]$, where ϵ is the optimism level.

To estimate λ_{tn_i} , we train a VAE $G_i^2 = \{E_i^2(\mu^2, \sigma^2 | s, a, s'), D_i^2(a | s, s', z^2 \sim (\mu^2, \sigma^2))\}$. Since the latent variable of VAE follows the Gaussian distribution, we use the mean as the encoding of the input and estimate the probability density functions: $\rho_i(s, a) \approx \rho_{\mathcal{N}(0,1)}(\mu_i^1)$ and $\rho_i(s, a, s') \approx \rho_{\mathcal{N}(0,1)}(\mu_i^2)$, where $\rho_{\mathcal{N}(0,1)}$ is the density of unit Gaussian distribution. The conditional density is

$\rho_i(s' | a, s) \approx \rho_{\mathcal{N}(0,1)}(\mu_i^2) / \rho_{\mathcal{N}(0,1)}(\mu_i^1)$ and the transition probability is $P_{\mathcal{B}_i}(s' | s, a_i) \approx \int_{s' - \frac{1}{2}\delta_S}^{s' + \frac{1}{2}\delta_S} \rho_i(s' | s, a) ds' \approx \rho_i(s' | s, a) \|\delta_S\|$ when $\|\delta_S\|$ is a small constant. Approximately, we have

$$\lambda_{tn_i} = \frac{\rho_{\mathcal{N}(0,1)}(\mu_i^1)}{\rho_{\mathcal{N}(0,1)}(\mu_i^2)},$$

and the constant $\|\delta_S\|$ is considered in z_i . In practice, we find that λ_{tn_i} falls into the region $[0.2, 1.4]$ for almost all samples. For completeness, we summarize the training of MABCQ in Algorithm 1.

Table 3: Transition probabilities and expected returns calculated in the dataset using only λ_{vd} .

	action	transition	expected return
Agent 1	a_1	$p(1 a_1) = 0.12$	4.52
		$p(5 a_1) = 0.88$	
	a_2	$p(6 a_2) = 0.8$	5
		$p(1 a_2) = 0.2$	
Agent 2	a_1	$p(1 a_1) = 0.4$	4
		$p(6 a_1) = 0.6$	
	a_2	$p(5 a_2) = 0.95$	4.8
		$p(1 a_2) = 0.05$	

Table 4: Transition probabilities and expected returns calculated in the dataset using λ_{tn} and λ_{vd} .

	action	transition	expected return
Agent 1	a_1	$p(1 a_1) = 0.17$	4.33
		$p(5 a_1) = 0.83$	
	a_2	$p(6 a_2) = 0.86$	5.29
		$p(1 a_2) = 0.14$	
Agent 2	a_1	$p(1 a_1) = 0.14$	5.29
		$p(6 a_1) = 0.86$	
	a_2	$p(5 a_2) = 0.83$	4.33
		$p(1 a_2) = 0.17$	

4 Related Work

4.1 Off-policy MARL

Many off-policy MARL methods have been proposed for learning to solve cooperative tasks in an online manner. Policy-based methods [14, 7, 36, 25, 28] extend actor-critic into multi-agent cases. Value factorization methods [26, 21, 24, 29] decompose the joint value function into individual value functions. All these methods follow CTDE, where the information of all agents can be accessed in a centralized way during training. Unlike these studies, we consider decentralized settings where global information is not available. For decentralized learning, the key challenge is the obsolete experiences in the replay buffer, which is considered in Fingerprints [2], Lenient-DQN [18], and concurrent experience replay [17]. However, these methods require additional information, e.g., training iteration number and exploration rate, which are often not provided by the offline dataset.

Table 5: Normalized scores of MABCQ and the baselines.

	MABCQ	BCQ w/ λ_{vd}	BCQ w/ λ_{tn}	BCQ	DDPG	Behavior
HalfCheetah	17.6 \pm 3.3	13.3 \pm 4.8	13.4 \pm 6.5	13.4 \pm 5.2	-3.1 \pm 3.6	11.3 \pm 2.8
Walker	54.4 \pm 5.6	50.1 \pm 11.0	41.2 \pm 17.8	28.8 \pm 14.4	1.7 \pm 0.9	10.0 \pm 0.8
Hopper	43.1 \pm 14.2	34.1 \pm 8.2	19.8 \pm 8.7	18.0 \pm 3.4	9.0 \pm 15.3	10.7 \pm 2.3
Ant	60.5 \pm 3.6	59.7 \pm 4.9	62.9 \pm 2.1	51.5 \pm 12.7	-48.4 \pm 1.3	19.3 \pm 4.5

4.2 Offline RL

Offline RL requires the agent to learn from a fixed batch of data consisting of single-step transitions, without exploration. Most offline RL methods consider the out-of-distribution action [11] as the fundamental challenge, which is the main cause of the extrapolation error [5] in value estimate in the single-agent environment. To minimize the extrapolation error, some recent methods introduce constraints to enforce the learned policy to be close to the behavior policy, which can be direct action constraint [5], kernel MMD [9], Wasserstein distance [30], KL divergence [20], or l_2 distance [4, 19]. Some methods train a Q-function pessimistic to out-of-distribution actions to avoid overestimation by adding a reward penalty quantified by the learned environment model [34], by minimizing the Q-values of out-of-distribution actions [10, 33], by weighting the update of Q-function via Monte Carlo dropout [31], or by explicitly assigning and training pseudo Q-values for out-of-distribution actions [15]. Our framework can be built on these methods.

MAICQ [32] studies offline MARL in the CTDE setting, which requires the joint actions of all agents in the dataset and cannot be applied to decentralized settings where datasets contain only individual actions. *All the methods aforementioned do not consider the extrapolation error introduced by the transition bias, which is a fatal problem in offline decentralized MARL.*

5 Experiments

We evaluate our framework in both fully and partially observable tasks. In each task, we build offline dataset \mathcal{B}_i for each agent i , which *does not contain actions of other agents*. We will give the details about the collection of each offline dataset. Our method and baselines have the same neural network architectures and hyperparameters, which are available in Appendix. All the models are trained for five runs with different random seeds, and the results are presented in terms of mean and std.

5.1 Matrix Game

We perform MABCQ on the matrix game in Table 1. As shown in Table 3, if we only use λ_{vd} without considering *transition normalization*, since the transition probabilities of high-value next states have been increased, for agent 1 the value of a_2 becomes higher than that of a_1 . However, due to the unbalanced action distribution of agent 1, the initial transition probabilities of agent 2 are extremely unbalanced. With λ_{vd} , agent 2 still underestimates the value of a_1 and learns the action a_2 . The agents arrive at the joint action (a_2, a_2) , which is a worse solution than the initial one (Table 2). Further, by normalizing the transition probabilities by λ_{tn} , the agents can learn the optimal solution (a_2, a_1) and build the consensus about the values of learned actions, as shown in Table 4.

Table 6: Mean difference in value estimates among agents during training. It is shown that *transition normalization* indeed reduces the difference in value estimates.

	value difference	MABCQ	BCQ w/ λ_{vd}
HalfCheetah		44.4 \pm 3.4	411.7 \pm 72.4
Walker		28.2 \pm 2.8	38.7 \pm 6.9
Hopper		24.2 \pm 0.8	25.4 \pm 1.3
Ant		60.8 \pm 2.9	67.0 \pm 3.1

Table 7: Extrapolation errors. It is shown that our framework can decrease the extrapolation error.

	extrapolation error	MABCQ	BCQ
HalfCheetah		98.4 \pm 31.3	97.2 \pm 29.1
Walker		55.0 \pm 9.6	91.5 \pm 35.4
Hopper		28.1 \pm 3.4	65.8 \pm 6.4
Ant		180.2 \pm 22.2	231.3 \pm 47

5.2 Multi-Agent Mujoco

To evaluate MABCQ in high-dimensional complex environments, we adopt multi-agent mujoco [1], where each agent independently controls one or some joints of the robot and can get the state [8] and reward of the robot. The task illustration and the collection of offline datasets are given in Appendix.

Baselines. We compare MABCQ against

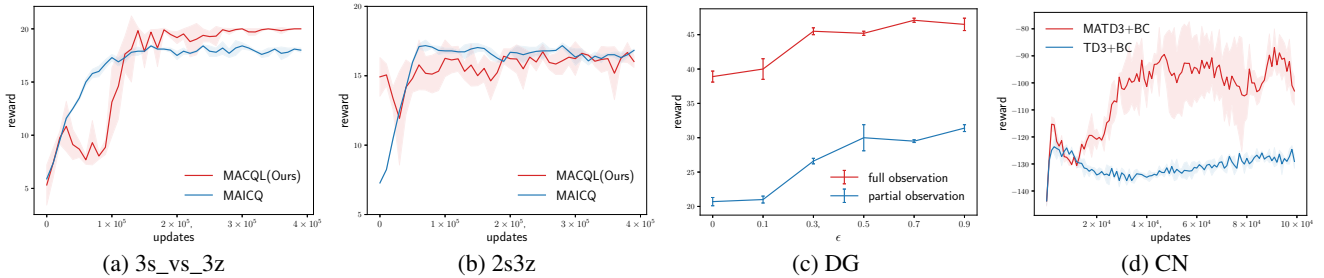
- BCQ w/ λ_{vd} . Using λ_{vd} alone on BCQ.
- BCQ w/ λ_{tn} . Using λ_{tn} alone on BCQ.
- BCQ. Removing both λ_{tn} and λ_{vd} from MABCQ.
- DDPG [13]. Each agent i is trained using independent DDPG on the offline \mathcal{B}_i without action constraint and transition probability modification.
- Behavior. Each agent i takes the action generated from the VAE G_i^1 .

Ablation. Table 5 shows the normalized scores [3] of all the methods in the four tasks. Without action constraint, DDPG severely suffers from the large extrapolation error and hardly learns. BCQ outperforms the behavior policies but only arrives at mediocre performance. Using λ_{tn} alone does not always improve the performance, *e.g.*, in HalfCheetah and Hopper. This is because λ_{tn} makes transition probabilities be uniform, which can be far from the ones in execution, leading to large extrapolation errors. In Ant, BCQ w/ λ_{tn} outperforms BCQ, which is attributed to the value consensus built by the normalized transition probabilities. By optimistically increasing the transition probabilities of high-value next states, λ_{vd} mitigates the underestimation of potential good actions and thus boosts the performance. MABCQ combines the advantages of value deviation and transition normalization and outperforms other baselines.

Consensus on value estimates. To verify that *transition normalization* can decrease the difference in value estimates among agents,

Table 8: Rewards on SMAC datasets.

		MABCQ	BCQ	MACQL	CQL
random	3m	4.0 \pm 1.1	0 \pm 0	13.5 \pm 1.5	9.3 \pm 3.0
	8m	4.5 \pm 0.9	0.8 \pm 0.2	8.2 \pm 1.1	6.6 \pm 1.2
	3s_vs_3z	8.1 \pm 0.7	0 \pm 0	10.2 \pm 0.8	10.1 \pm 0.4
	3s_vs_4z	4.1 \pm 1.5	0 \pm 0	5.7 \pm 0.6	6.7 \pm 0.2
medium	3m	8.9 \pm 1.3	7.8 \pm 0.4	15.1 \pm 1.8	13.8 \pm 1.4
	8m	7.6 \pm 0.8	4.5 \pm 1.2	14.5 \pm 1.5	12.4 \pm 0.6
	3s_vs_3z	8.7 \pm 1.1	3.9 \pm 0.6	9.3 \pm 0.9	8.9 \pm 0.6
	3s_vs_4z	4.3 \pm 0.5	0 \pm 0	6.1 \pm 0.7	6.8 \pm 1.8
replay	3m	13.2 \pm 0.2	12.7 \pm 0.7	13.8 \pm 0.4	13.5 \pm 0.6
	8m	15.2 \pm 1.0	14.3 \pm 0.9	17.9 \pm 0.4	16.3 \pm 0.4
	3s_vs_3z	19.4 \pm 0.4	19.8 \pm 0.3	20.0 \pm 0.0	20.0 \pm 0.0
	3s_vs_4z	5.3 \pm 0.6	5.3 \pm 0.9	5.9 \pm 0.3	5.2 \pm 0.7
expert	3m	18.8 \pm 0.7	18.3 \pm 1.1	18.9 \pm 0.5	19.1 \pm 0.5
	8m	17.0 \pm 0.8	17.5 \pm 1.1	18.5 \pm 1.2	18.3 \pm 1.1
	3s_vs_3z	19.1 \pm 0.5	19.0 \pm 0.6	19.2 \pm 0.6	19.1 \pm 0.6
	3s_vs_4z	5.6 \pm 0.9	5.4 \pm 1.1	6.8 \pm 0.7	6.5 \pm 0.9

**Figure 1:** (a) and (b): Comparison with MAICQ on two SMAC datasets [32]. (c): Performance of MATD3+BC with different ϵ on DG datasets. (d): Learning curves on CN datasets.

we uniformly sample a subset from the union of all agents’ states and calculate the difference in value estimates, $\max_i V_i^* - \min_i V_i^*$, on this subset. The mean differences during training are illustrated in Table 6. The $\max_i V_i^* - \min_i V_i^*$ of MABCQ is indeed lower than that of BCQ w/ λ_{vd} . If there is a consensus among agents about which states are high-value, the agents will select the actions that most likely lead to the common high-value states. This promotes the coordination of policies and helps MABCQ outperform BCQ w/ λ_{vd} .

Extrapolation error. In Table 7, we present the extrapolation errors of MABCQ and BCQ, measured by $|\frac{1}{N} \sum_i Q_i(s, a_i) - G|$, where G is the true return evaluated by Monte Carlo. Although MABCQ greatly outperforms BCQ (*i.e.*, much higher return), it still achieves much smaller extrapolation errors than BCQ in Walker, Hopper, and Ant. This empirically verifies the claim that our method can decrease the extrapolation error.

5.3 SMAC

We also investigate the proposed framework on SMAC [22] tasks, including 3m, 8m, 3s_vs_3z, and 3s_vs_4z. We build random datasets that are generated by uniform policies, medium datasets that are generated by mixed medium and uniform policies, replay datasets that are collected in the training process of QMIX [21], and expert datasets that are generated by expert policies trained by QMIX. Each dataset contains 1×10^4 episodes. We also build our framework on CQL [10], as MACQL. As shown in Table 8, MABCQ and MACQL achieve great performance improvement compared with the base-

lines, especially in random and medium datasets, where the transition dynamics in execution are much different from the ones in offline datasets. In expert datasets, since the behavior policies are much deterministic, offline RL methods avoid selecting out-of-distribution actions and thus degenerate to behavior cloning. Therefore, all the methods perform similarly.

Offline CTDE settings. Although offline CTDE method [32] does not fit our offline decentralized setting, our framework can work in offline CTDE datasets. To verify, we select two replay datasets (jointly collected) in MAICQ [32], split them into individual datasets, and test MACQL on them. As shown in Figures 1a and 1b, our decentralized method can obtain competitive performance compared with the centralized method, MAICQ.

Table 9: Rewards on DG datasets.

	MATD3+BC	TD3+BC
full observation	46.5 \pm 0.9	38.9 \pm 0.8
partial observation	31.4 \pm 0.6	20.7 \pm 0.6

5.4 MPE

We additionally evaluate our framework in an MPE-based [14] Differential Game (DG), where the transition bias greatly affects the performance. Two agents can move in the range $[-1, 1]$. The action is the speed, which is in the range $[-0.1, 0.1]$. Define $l = \sqrt{x_1^2 + x_2^2}$,


Table 10: Normalized scores on D4RL MuJoCo datasets.

		MATD3+BC	TD3+BC	TD3+BC (single)
random	halfcheetah	14.3 \pm 2.9	12.9 \pm 2.6	10.2 \pm 1.3
	hopper	10.3 \pm 0.6	10.4 \pm 0.8	11.0 \pm 0.1
	walker2d	4.5 \pm 3.4	3.7 \pm 3.0	1.4 \pm 1.6
medium	halfcheetah	41.5 \pm 2.4	40.4 \pm 2.4	42.8 \pm 0.3
	hopper	97.1 \pm 1.7	98.7 \pm 1.2	99.5 \pm 1.0
	walker2d	82.0 \pm 5.1	72.3 \pm 4.1	79.7 \pm 1.8
replay	halfcheetah	40.0 \pm 3.2	39.3 \pm 2.8	43.3 \pm 0.5
	hopper	28.4 \pm 4.1	25.5 \pm 3.3	31.4 \pm 3.0
	walker2d	21.0 \pm 2.9	21.3 \pm 1.5	25.2 \pm 5.1
medium-expert	halfcheetah	96.2 \pm 5.5	95.3 \pm 6.5	97.9 \pm 4.4
	hopper	112.0 \pm 0.8	112.3 \pm 0.9	112.2 \pm 0.2
	walker2d	88.9 \pm 17.4	82.0 \pm 10.8	105.7 \pm 2.7

where x_1 and x_2 are the positions of the two agents, respectively. The shared reward is set as

$$r = \begin{cases} 0.5 \times (\cos(15 \times l) + 1) & \text{if } l < 0.2 \\ 0 & \text{if } 0.2 \leq l \leq 0.6 \\ 0.5 \times (l - 0.6)^2 & \text{if } l > 0.6. \end{cases}$$

The visualization of reward function is shown in Figure 2.

**Figure 2:** Visualization of reward function in Differential Game.

Partial observation vs full observation. The offline datasets are collected by uniform random policies, containing 1×10^6 transitions. In the full observation setting, the dataset of each agent contains the positions of both agents. In the partial observation setting, the dataset of each agent only contains its own position. In both settings, the datasets do not contain the actions of the other agent. We add λ_{vd} and λ_{tn} to TD3+BC [4], as MATD3+BC. As shown in Table 9, MATD3+BC obtains the substantial improvement in both full and partial observation settings.

Hyperparameter ϵ . The optimism level ϵ controls the strength of value deviation. If ϵ is too small, value deviation has weak effects on the objective function. On the other hand, if ϵ is too large, the agent will be overoptimistic about other agents’ learned policies. Figure 1c shows that the performance of MATD3+BC with different ϵ , which verifies our framework is robust to ϵ .

We test MATD3+BC on Cooperative Navigation in MPE, where 4 agents learn to cover 4 landmarks. The reward is $-\text{sum}(\text{distance}_j)$, where distance_j is the distance from landmark j to the closest agent. The offline datasets are collected by uniform random policies, containing 1×10^6 transitions. Figure 1d shows our framework significantly outperforms the baseline.

5.5 Additional Results

We also split the D4RL [3] Mujoco datasets into decentralized multi-agent datasets, and test MATD3+BC on them. The results are summarized in Table 10. We find the results of decentralized methods, where the joints are controlled by different agents, are very close to the results of single-agent method, TD3+BC (single), where a single agent controls all joints, which could be seen as the “upper-bound” of the decentralized methods. That is the reason that our method does not bring significant improvement in these tasks. However, MATD3+BC still outperforms TD3+BC on several tasks, *e.g.*, halfcheetah-random and walker2d-medium.

Table 11: Average time taken by one update.

MABCQ	BCQ	MATD3+BC	TD3+BC
18 ms	10 ms	4 ms	3 ms

To demonstrate the computation efficiency of our method, in Table 11, we record the average time taken by one update in Halfcheetah. The experiments are carried out on Intel i7-8700 CPU and NVIDIA GTX 1080Ti GPU. Since λ_{vd} and λ_{tn} could be calculated from the sampled experience without actually going through all next states, our framework additionally needs only two forward passes for computing λ_{vd} and λ_{tn} in the update. Since the value computation is very efficient in TD3+BC, our framework is also efficient on it.

6 Conclusion

We propose a framework for offline decentralized multi-agent reinforcement learning, to overcome the mismatch between transition dynamics. The framework can be instantiated on many offline RL methods. Theoretically, we show that under the purposely controlled non-stationary transition dynamics, offline decentralized Q-learning converges to a unique fixed point. Empirically, the framework outperforms the baselines in a variety of multi-agent offline datasets.

Acknowledgements

This work was supported by NSF China under grants 62250068 and 61872009. The authors would like to thank the anonymous reviewers for their valuable comments.

References

- [1] Christian Schroeder de Witt, Bei Peng, Pierre-Alexandre Kamienny, Philip Torr, Wendelin Böhmer, and Shimon Whiteson, 'Deep Multi-Agent Reinforcement Learning for Decentralized Continuous Cooperative Control', *arXiv preprint arXiv:2003.06709*, (2020).
- [2] Jakob Foerster, Nantas Nardelli, Gregory Farquhar, Triantafyllos Afouras, Philip HS Torr, Pushmeet Kohli, and Shimon Whiteson, 'Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning', in *International Conference on Machine Learning (ICML)*, (2017).
- [3] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine, 'D4rl: Datasets for Deep Data-Driven Reinforcement Learning', *arXiv preprint arXiv:2004.07219*, (2020).
- [4] Scott Fujimoto and Shixiang Shane Gu, 'A Minimalist Approach To Offline Reinforcement Learning', in *Thirty-Fifth Conference on Neural Information Processing Systems (NeurIPS)*, (2021).
- [5] Scott Fujimoto, David Meger, and Doina Precup, 'Off-Policy Deep Reinforcement Learning Without Exploration', in *International Conference on Machine Learning (ICML)*, (2019).
- [6] Julian Ibarz, Jie Tan, Chelsea Finn, Mrinal Kalakrishnan, Peter Pastor, and Sergey Levine, 'How To Train Your Robot with Deep Reinforcement Learning: Lessons We Have Learned', *The International Journal of Robotics Research*, **40**(4-5), 698–721, (2021).
- [7] Shariq Iqbal and Fei Sha, 'Actor-Attention-Critic for Multi-Agent Reinforcement Learning', in *International Conference on Machine Learning (ICML)*, (2019).
- [8] Jakub Grudzien Kuba, Ruiqing Chen, Munning Wen, Ying Wen, Fanglei Sun, Jun Wang, and Yaodong Yang, 'Trust Region Policy Optimization in Multi-Agent Reinforcement Learning', *International Conference on Learning Representations (ICLR)*, (2022).
- [9] Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine, 'Stabilizing Off-Policy Q-Learning Via Bootstrapping Error Reduction', in *Advances in Neural Information Processing Systems (NeurIPS)*, (2019).
- [10] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine, 'Conservative Q-Learning for Offline Reinforcement Learning', *Neural Information Processing Systems (NeurIPS)*, (2020).
- [11] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu, 'Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems', *arXiv preprint arXiv:2005.01643*, (2020).
- [12] Lihong Li, Wei Chu, John Langford, and Robert E Schapire, 'A Contextual-Bandit Approach To Personalized News Article Recommendation', in *International Conference on World Wide Web (WWW)*, pp. 661–670, (2010).
- [13] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra, 'Continuous Control with Deep Reinforcement Learning', in *International Conference on Learning Representations (ICLR)*, (2016).
- [14] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch, 'Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments', in *Advances in Neural Information Processing Systems (NeurIPS)*, (2017).
- [15] Jiafei Lyu, Xiaoteng Ma, Xiu Li, and Zongqing Lu, 'Mildly Conservative Q-learning for Offline Reinforcement Learning', *Neural Information Processing Systems (NeurIPS)*, (2022).
- [16] Frans A Oliehoek and Christopher Amato, *A Concise Introduction To Decentralized POMDPs*, Springer, 2016.
- [17] Shayegan Omidshafiei, Jason Papis, Christopher Amato, Jonathan P How, and John Vian, 'Deep Decentralized Multi-Task Multi-Agent Reinforcement Learning Under Partial Observability', in *International Conference on Machine Learning (ICML)*, (2017).
- [18] Gregory Palmer, Karl Tuyls, Daan Bloembergen, and Rahul Savani, 'Lenient Multi-Agent Deep Reinforcement Learning', in *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, (2018).
- [19] Ling Pan, Longbo Huang, Tengyu Ma, and Huazhe Xu, 'Plan Better Amid Conservatism: Offline Multi-Agent Reinforcement Learning with Actor Rectification', *arXiv preprint arXiv:2111.11188*, (2021).
- [20] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine, 'Advantage-Weighted Regression: Simple and Scalable Off-Policy Reinforcement Learning', *arXiv preprint arXiv:1910.00177*, (2019).
- [21] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson, 'QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning', in *International Conference on Machine Learning (ICML)*, (2018).
- [22] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob Foerster, and Shimon Whiteson, 'The StarCraft Multi-Agent Challenge', *CoRR*, **abs/1902.04043**, (2019).
- [23] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz, 'Trust Region Policy Optimization', in *International Conference on Machine Learning (ICML)*, (2015).
- [24] Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi, 'QTRAN: Learning To Factorize with Transformation for Cooperative Multi-Agent Reinforcement Learning', in *International Conference on Machine Learning (ICML)*, (2019).
- [25] Kefan Su and Zongqing Lu, 'Divergence-Regularized Multi-Agent Actor-Critic', in *International Conference on Machine Learning (ICML)*, (2022).
- [26] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al., 'Value-Decomposition Networks for Cooperative Multi-Agent Learning Based on Team Reward', in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, (2018).
- [27] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al., 'Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning', *Nature*, **575**(7782), 350–354, (2019).
- [28] Jiangxing Wang, Deheng Ye, and Zongqing Lu, 'More Centralized Training, Still Decentralized Execution: Multi-Agent Conditional Policy Factorization', in *International Conference on Learning Representations (ICLR)*, (2023).
- [29] Jianhao Wang, Zhizhou Ren, Terry Liu, Yang Yu, and Chongjie Zhang, 'Qplex: Duplex Dueling Multi-Agent Q-Learning', in *International Conference on Learning Representations (ICLR)*, (2021).
- [30] Yifan Wu, George Tucker, and Ofir Nachum, 'Behavior Regularized Offline Reinforcement Learning', *arXiv preprint arXiv:1911.11361*, (2019).
- [31] Yue Wu, Shuangfei Zhai, Nitish Srivastava, Joshua Susskind, Jian Zhang, Ruslan Salakhutdinov, and Hanlin Goh, 'Uncertainty Weighted Actor-Critic for Offline Reinforcement Learning', *arXiv preprint arXiv:2105.08140*, (2021).
- [32] Yiqin Yang, Xiaoteng Ma, Chenghao Li, Zewu Zheng, Qiyuan Zhang, Gao Huang, Jun Yang, and Qianchuan Zhao, 'Believe What You See: Implicit Constraint Approach for Offline Multi-Agent Reinforcement Learning', in *Neural Information Processing Systems (NeurIPS)*, (2021).
- [33] Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn, 'Combo: Conservative Offline Model-Based Policy Optimization', *arXiv preprint arXiv:2102.08363*, (2021).
- [34] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma, 'MOPO: Model-Based Offline Policy Optimization', *Advances in Neural Information Processing Systems (NeurIPS)*, (2020).
- [35] Chi Zhang, Sanmukh Rao Kuppannagari, and Viktor Prasanna, BRAC+: Going Deeper with Behavior Regularized Offline Reinforcement Learning, 2021.
- [36] Tianhao Zhang, Yueheng Li, Chen Wang, Guangming Xie, and Zongqing Lu, 'FOP: Factorizing Optimal Joint Policy of Maximum-Entropy Multi-Agent Reinforcement Learning', in *International Conference on Machine Learning (ICML)*, (2021).