

# ME-IGM: Individual-Global-Max in Maximum Entropy Multi-Agent Reinforcement Learning

Wen-Tse Chen  
Carnegie Mellon University  
Pittsburgh, United States  
wentsec@andrew.cmu.edu

Yuxuan Li  
Zhejiang University  
Hangzhou, China  
yuxuanli04@zju.edu.cn

Shiyu Huang  
XPeng Inc.  
Beijing, China  
huangsy16@xiaopeng.com

Jiayu Chen\*  
The University of Hong Kong  
INFIFORCE Intelligent Tech. Co., Ltd.  
Hong Kong, China  
jiayuc@hku.hk

Jeff Schneider\*  
Carnegie Mellon University  
Pittsburgh, United States  
jeff4@andrew.cmu.edu

## ABSTRACT

Multi-agent credit assignment is a fundamental challenge for cooperative multi-agent reinforcement learning (MARL), where a team of agents learn from shared reward signals. The Individual-Global-Max (IGM) condition is a widely used principle for multi-agent credit assignment, requiring that the joint action determined by individual Q-functions maximizes the global Q-value. Meanwhile, the principle of maximum entropy has been leveraged to enhance exploration in MARL. However, we identify a critical limitation in existing maximum entropy MARL methods: a misalignment arises between local policies and the joint policy that maximizes the global Q-value, leading to violations of the IGM condition. To address this misalignment, we propose an order-preserving transformation. Building on it, we introduce ME-IGM, a novel maximum entropy MARL algorithm compatible with any credit assignment mechanism that satisfies the IGM condition while enjoying the benefits of maximum entropy exploration. We empirically evaluate two variants of ME-IGM: ME-QMIX and ME-QPLEX, in non-monotonic matrix games, and demonstrate their state-of-the-art performance across 17 scenarios in SMAC-v2 and Overcooked.

## KEYWORDS

Multi Agent Reinforcement Learning; Maximum Entropy Reinforcement Learning; Centralized Training with Decentralized Execution

### ACM Reference Format:

Wen-Tse Chen, Yuxuan Li, Shiyu Huang, Jiayu Chen, and Jeff Schneider. 2026. ME-IGM: Individual-Global-Max in Maximum Entropy Multi-Agent Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 19 pages. <https://doi.org/10.65109/GYYC3346>

\*Co-supervision.



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

## 1 INTRODUCTION

Collaborative multi-agent tasks, where a team of agents works together to complete a task and receives joint rewards, are inherently challenging for reinforcement learning (RL) agents. This difficulty arises due to the absence of individual reward signals for each agent and the significantly higher exploration requirements compared to single-agent scenarios. The challenge is further exacerbated by the exponential growth of the joint action space as the number of agents increases. In this work, we propose leveraging maximum entropy RL to promote exploration and adopting the individual-global-max (IGM) credit assignment mechanism [22] to effectively distribute joint rewards among the agents.

Maximum entropy RL [14], recognized for promoting exploration [9], smoothing the objective landscape [1], and improving robustness [6], has been extensively applied in single-agent RL. However, as illustrated in Figure 1, naively extending maximum entropy RL to multi-agent settings under the centralized training and decentralized execution (CTDE) paradigm leads to uniform credit assignment, which limits the ability of multi-agent reinforcement learning (MARL) agents to learn from joint rewards.

On the other hand, numerous studies have explored the development of effective credit assignment mechanisms. Among them, the IGM condition has demonstrated strong empirical performance by aligning the maximum local Q-values with the maximum global Q-value. Previous works have primarily applied the IGM condition in deterministic policy settings, where policies are derived using the argmax of local Q-values. However, in maximum entropy MARL methods, stochastic local policies are learned, and misalignment can occur between local policies and the maximum global Q-values. **That is, selecting joint actions based on the highest logits of local policies does not necessarily lead to maximizing the global Q-value.** This limitation highlights a critical challenge in achieving effective credit assignment within the maximum entropy MARL framework. This challenge, as we will demonstrate, is not unique to IGM-based methods but also represents a broader obstacle for actor-critic-based maximum entropy MARL algorithms, such as FOP [28], hindering their ability to achieve optimal coordination.

We propose ME-IGM, the first CTDE maximum entropy MARL method that is compatible with the IGM condition. In particular, to address the misalignment between stochastic local policies and the maximum global Q-values, we introduce an order-preserving

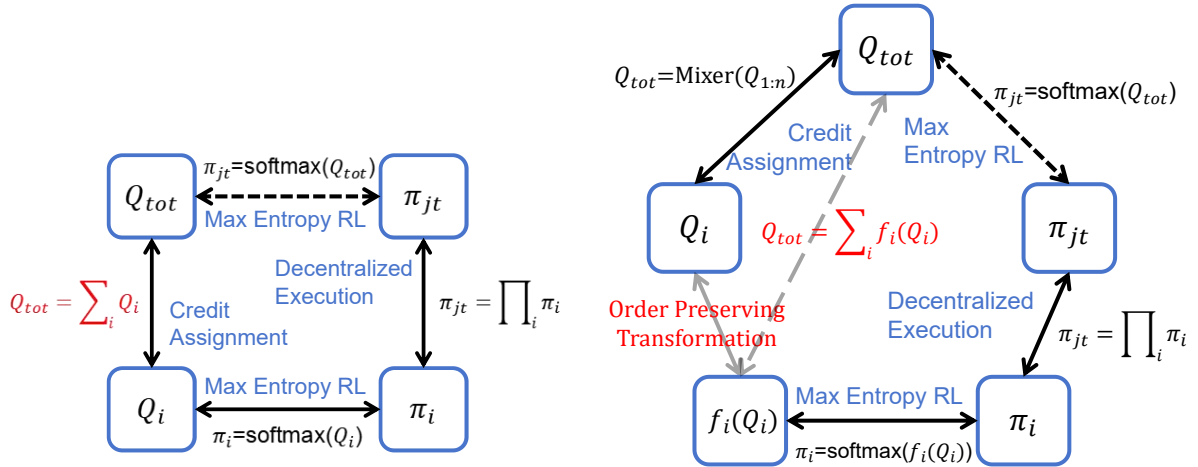


Figure 1: The figure illustrates the improvement of our approach compared to existing maximum entropy MARL methods. The left figure shows a straightforward approach to applying maximum entropy MARL in the CTDE context, where blue texts represent the desired objectives, and black texts indicate corresponding constraints. It reveals that existing maximum entropy MARL methods under the CTDE framework implicitly constrain the global Q-value to be the sum of local Q-values (as in VDN), significantly limiting the expressiveness of the critic network. The right figure depicts the improvements in ME-IGM. ME-IGM first applies any credit assignment mechanism that satisfies the IGM condition, such as QMIX and QPLEX, to obtain  $Q_i$  for each agent  $i$ . Given the meaningful order of local Q-values (for different actions), an order-preserving transformation  $f_i$  converts  $Q_i$  to  $f_i(Q_i)$  as the policy logits. This transformation is optimized using a loss function that minimizes the expected difference between  $\sum_i f_i(Q_i)$  and  $Q_{tot}$ , which guarantees monotonic policy improvement in maximum entropy MARL.

transformation (OPT) that maps local Q-values to policy logits while preserving their relative order. This ensures that selecting actions with the highest logits collectively leads to a joint action that maximizes the global Q-value. Additionally, the OPT is trained with a theoretically grounded objective, guaranteeing monotonic policy improvement in CTDE MARL.

Our main contributions are summarized as follows:

- We are the first to identify and empirically demonstrate a critical issue of existing maximum entropy MARL methods: the misalignment between stochastic local policies and the maximum global Q-values.
- We introduce ME-IGM, the first CTDE maximum entropy MARL method that fully complies with the IGM condition. The core of ME-IGM is an order-preserving transformation (OPT), which maps local Q-values to policy logits while preserving their relative order. This ensures that selecting actions with the highest logits collectively leads to a joint action that maximizes the global Q-value. OPT can be seamlessly integrated into mainstream value decomposition MARL methods for a maximum entropy extension to enhance exploration. Additionally, we propose a theoretically grounded and straightforward objective for training these transformation operators.
- We validate our algorithm on matrix games, SMAC-v2 [5], and Overcooked [3], where our method attains state-of-the-art results.<sup>1 2</sup>

<sup>1</sup>A full version of this paper with appendices is available on [arXiv]

<sup>2</sup>Our implementation is publicly available on [GitHub]

## 2 BACKGROUND

### 2.1 Multi-Agent Reinforcement Learning

Cooperative MARL can be formalized as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) [18]. Formally, a Dec-POMDP is represented as a tuple  $(\mathcal{A}, S, U, T, r, O, G, \gamma)$ , where  $\mathcal{A} \equiv \{1, \dots, n\}$  is the set of  $n$  agents.  $S, U, O$ , and  $\gamma$  are state space, action space, observation space, and discount factor, respectively.

During each discrete time step  $t$ , each agent  $i \in \mathcal{A}$  chooses an action  $u^i \in U$ , resulting in a collective joint action  $\mathbf{u} \in U \equiv U^n$ . The function  $r(s, \mathbf{u})$  defines the immediate reward for all agents when the collective action  $\mathbf{u}$  is taken in the state  $s$ .  $\mathcal{P}(s, \mathbf{u}, s') : S \times U \times S \rightarrow [0, \infty)$  is the state-transition function, which defines the probability density of the succeeding state  $s'$  after taking action  $\mathbf{u}$  in state  $s$ . In a Dec-POMDP, each agent receives only partial observations  $o \in O$  according to the observation function  $G(s, i) : S \times \mathcal{A} \rightarrow O$ . For simplicity, if a function depends on both  $o^i$  and  $s$ , we will disregard  $o^i$ . Each agent uses a policy  $\pi_i(u^i|o^i)$  to produce its action  $u^i$ . Note that the policy should be conditioned on the observation history  $o_{1:t}^i$ . For simplicity, we refer to  $o_{1:t}^i$  as  $o_t^i$  or  $o^i$  in the whole paper. The joint policy is denoted as  $\pi_{jt}$ .  $\rho_\pi(s_t, \mathbf{u}_t)$  is used to represent the state-action marginals of the trajectory distribution induced by a policy  $\pi$ .

### 2.2 Individual-Global-Max Condition

Individual-Global-Max (IGM) [22] is a commonly used credit assignment method in value-decomposition-based cooperative MARL and is defined as follows:

*Definition 2.1. Individual-Global-Max (IGM)* For joint q-function  $Q_{tot}$ , if there exist individual q-functions  $[Q_i]_{i=1}^n$  such that the following holds:

$$\operatorname{argmax}_{\mathbf{u}_t} Q_{tot}(\mathbf{u}_t, s_t) = \begin{pmatrix} \operatorname{argmax}_{u_t^1} Q_1(u_t^1 | o_t^1) \\ \vdots \\ \operatorname{argmax}_{u_t^n} Q_n(u_t^n | o_t^n) \end{pmatrix}. \quad (1)$$

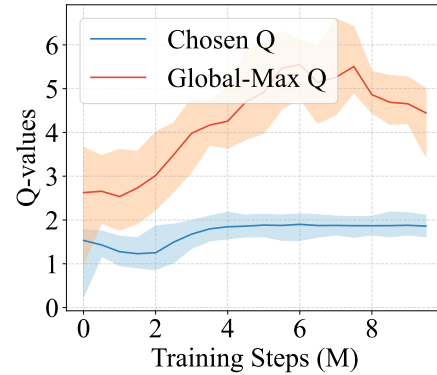
Then, we say that  $[Q_i]_{i=1}^n$  satisfy **IGM** for  $Q_{tot}$  under  $s_t$ .

### 2.3 Related Works

**Multi-Agent RL.** In the CTDE framework, MARL primarily involves two common algorithmic approaches: policy gradient methods [7, 13, 15, 17, 27, 29] and value function decomposition methods [22, 25, 26]. Policy gradient methods first learn a centralized critic network, and then distill local policies using losses like the KL divergence. On the other hand, Value function decomposition addresses the credit assignment problem by decomposing the global value function into multiple local value functions. VDN [25] assumes that the global Q-function is the sum of local Q-functions. QMIX [19, 22] permits the mixer function to be any function with non-negative weights. QTRAN [24] optimizes an additional pair of inequality constraints to construct a loss function. QPLEX [26] use the dueling architecture to decompose the Q-function, achieving the same expressive power as QTRAN while being easier to optimize. The expressive power of the critic networks in these algorithms increases sequentially, but experimental results show that QMIX has better performance [11]. Therefore, in this work, QMIX is used as the mechanism for credit assignment.

**Maximum Entropy MARL.** Maximum entropy RL [14] is demonstrated to have advantages such as encouraging exploration [9] and increasing robustness [6] in single-agent RL scenarios. In maximum entropy MARL, most works utilize an actor-critic (AC) architecture and typically aim to minimize the KL divergence between the local actor and the centralized critic [8, 10, 21, 28].

**Comparison with Actor-Critic Approaches.** Most existing maximum entropy MARL methods (e.g., mSAC [21], FOP [28], and HASAC [16]) rely on the Actor-Critic framework. Our value-based ME-IGM differs from them in three aspects. First, **Order Preservation:** AC-based methods may suffer from a misalignment where the order of actions preferred by the local actor does not match the local Q-function due to approximation errors. In contrast, our method uses Order-Preserving Transformations (OPT) to strictly enforce this alignment, ensuring the IGM condition is met. Second, **Objective Function:** AC methods typically minimize KL divergence to distill policies. Our method minimizes the MSE loss between logits, which has been shown to be superior in distillation contexts [12]. Third, **Specific Baselines:** FOP [28] employs a credit assignment mechanism similar to QPLEX but relies on the Individual-Global-Optimal (IGO) assumption. While FOP argues that locally optimizing policies reduces the joint policy gap, it does not explicitly guarantee alignment. HASAC [16] avoids the IGO assumption through sequential optimization but ignores the misalignment problem between stochastic policies and Q-values. Our work identifies this misalignment as a critical bottleneck and



**Figure 2: Illustration of misalignment between local policies and the maximum global Q-value.** When naively combining the IGM condition with maximum entropy MARL, local policies often select suboptimal joint actions, leading to a lower Q-value, depicted by the blue curve. In contrast, the optimal joint action should achieve the global-max Q-value, shown as the orange curve.

proposes OPT to resolve it, thereby achieving state-of-the-art performance.

### 3 MISALIGNMENT BETWEEN LOCAL POLICIES AND THE MAXIMUM GLOBAL Q-VALUE

**The Misalignment Problem of Maximum Entropy MARL:** While the IGM condition defines the relationship between local Q-values and the global Q-value, what is often more important in practice is the alignment between local policies and the global Q-value. In other words, joint actions that maximize the global Q-value can be obtained by querying local policies. Previous work in maximum entropy MARL [8, 28] often decompose the global Q-value into individual local Q-values to satisfy the IGM condition. These local Q-values are then used to guide the training of local stochastic policies. Compared with value decomposition methods such as QMIX, they train stochastic policy networks in addition to the Q-functions, which can potentially improve multi-agent exploration in complex environments. **However, while the value decomposition ensures alignment between global and local Q-values via the IGM condition, there is typically no explicit constraints imposed between the local policies and the local Q-values during the policy training phase. This lack of coupling often leads to misalignment during testing, where the local policies fail to consistently select joint actions that maximize the global Q-value, defeating the purpose of enforcing the IGM condition.** Consequently, the learned policies may struggle to achieve optimal global coordination during execution.

**Empirical Demonstration:** To analyze the performance gap caused by the misalignment between local policies and the maximum global Q-value, we evaluate the difference between the maximum global Q-value and the global Q-value selected by local policies in the SMAC-v2 *zerg 5vs5* scenario, using a state-of-the-art Maximum

Entropy MARL algorithm [28]. Figure 2 shows the results, averaged over five random seeds with standard deviations included. We can observe that there is a significant Q-value gap, emphasizing the potential for improvement by developing algorithms that enable the selection of joint actions with the highest global Q-value through local stochastic policies.

## 4 ORDER-PRESERVING TRANSFORMATION

Here, we propose a novel order-preserving transformation (OPT) to solve the aforementioned misalignment problem between local policies and the maximum global Q-value.

### 4.1 Formulation of OPT

OPT is used to map local Q-values to the logits of local policies while preserving their relative order. Formally, the OPT is defined as:

$$\text{OPT}(x, s) = W^2 \sigma(W^1 x + b^1) + b^2 \quad (2)$$

Here,  $W^1_{(d_1 \times 1)}$ ,  $W^2_{(1 \times d_1)}$ ,  $b^1_{(d_1 \times 1)}$ , and  $b^2_{(1 \times 1)}$  are generated by a hyper-network<sup>3</sup> which takes the global state  $s$  as input. Consider  $x$  as a local Q-value of dimension  $(1 \times d_2)$ , then  $W^1 x + b^1$  is a set of hidden variables of size  $(d_1 \times d_2)$ . Notably, when all entries in  $W^1$  are non-negative, the relative order of elements in  $x$  is preserved in each of the  $d_1$  hidden variables. Next, a nonlinear activation function  $\sigma$ , such as ELU [4], is applied, followed by another order-preserving linear transformation using  $W^2$  and  $b^2$ . Consequently, the final output,  $\text{OPT}(x, s)$ , maintains both the dimension and the relative order of the entries in the input  $x$ .

$\text{OPT}(x, s)$  is used as the logits of the local policy corresponding to the local Q-value  $x$ . Through aforementioned order-preserving design, each agent can select its action corresponding to the highest local policy logit, which is also the maximum point of its local Q-function. The local Q functions can be learned with any value decomposition methods satisfying the IGM condition, to ensure that the selected local actions collectively compose the optimal joint action that maximize the global Q-function. Thus, the proposed OPT addresses the aforementioned misalignment issue commonly found in maximum entropy MARL.

### 4.2 Training Scheme of OPT

Here, we first formulate the ideal objective of Maximum Entropy MARL and its practical counterpart that admits an analytical-form solution. We then analyze the gap between these two objective functions, providing a theoretically sound objective design for Maximum Entropy MARL based on OPTs.

Maximum Entropy MARL trains the individual and global Q-functions as in value decomposition MARL methods, while training a policy network in a Maximum Entropy RL framework based on the global Q-values. Formally, the joint policy is trained as follows:

$$\min_{\pi'_{jt} \in \Pi} D_{KL} \left( \pi'_{jt}(\cdot | s_t) \left\| \frac{\exp(Q_{tot}^{\text{old}}(s_t, \cdot) / \alpha)}{Z^{\text{old}}(s_t)} \right. \right), \quad (3)$$

where  $Z^{\text{old}}(s_t) = \sum_{\mathbf{u}} \exp(Q_{tot}^{\text{old}}(s_t, \mathbf{u}) / \alpha)$  and  $\alpha$  is the temperature in the softmax distribution.

<sup>3</sup>For details of the hyper-network, please refer to the Appendix.

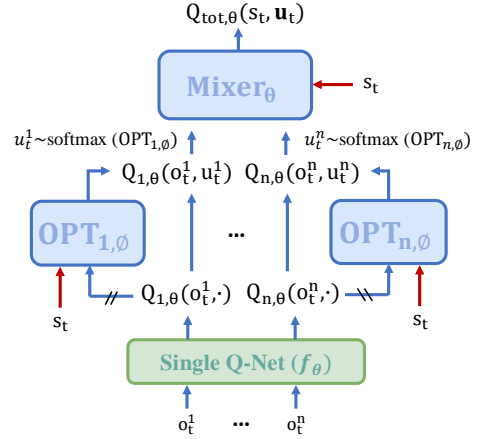


Figure 3: The overall pipeline of ME-IGM.

To simplify and decentralize the policy improvement process, we propose replacing  $Q_{tot}^{\text{old}}(s_t, \cdot)$  with  $\sum_i \text{OPT}_i(Q_i^{\text{old}}(o_t^i, \cdot), s_t)$ , which leads to:

$$\min_{\pi'_{jt} \in \Pi} D_{KL} \left( \pi'_{jt}(\cdot | s_t) \left\| \frac{\exp(\sum_i \frac{\text{OPT}_i(Q_i^{\text{old}}(o_t^i, \cdot), s_t)}{\alpha})}{Z'^{\text{old}}(s_t)} \right. \right), \quad (4)$$

where  $Z'^{\text{old}}(s_t)$  is the normalization factor. The joint policy  $\pi'_{jt}$  is implemented as a set of independent local policies for decentralized execution, and we have the analytical solution for the local policy of agent  $i$  as:

$$\pi_i^{\text{new}}(a | o_t^i) \propto \exp(\text{OPT}_i(Q_i^{\text{old}}(o_t^i, a), s_t) / \alpha). \quad (5)$$

We can see that the policy is defined based on the local Q-value and utilizes the OPT to ensure that the relative order of policy logits aligns with the local Q-values.

**There is a gap between  $\sum_i \text{OPT}_i(Q_i^{\text{old}}(o_t^i, \cdot), s_t)$  and  $Q_{tot}^{\text{old}}(s_t, \cdot)$ .** Also, we find that the policy improvement using Equation (4) no longer guarantees a monotonic increase in the global Q-value, unlike the approach based on Equation (3). Specifically, we have the following theorem:

**THEOREM 4.1.** *Denote the policy before the policy improvement as  $\pi_{\text{old}}$  and the policy achieving the optimality of the improvement step as  $\pi_{\text{new}}$ . Updating the policy with Equation (3) ensures monotonic improvement in the global Q-value. That is,  $Q_{tot}^{\text{old}}(s_t, \mathbf{u}_t) \leq Q_{tot}^{\text{new}}(s_t, \mathbf{u}_t), \forall s_t, \mathbf{u}_t$ .*

*In contrast, updating the policy with Equation (4) only ensures that  $Q_{tot}^{\text{old}}(s_t, \mathbf{u}_t) - \epsilon \leq Q_{tot}^{\text{new}}(s_t, \mathbf{u}_t), \forall s_t, \mathbf{u}_t$ . The gap  $\epsilon$  satisfies:*

$$\epsilon \leq \frac{\gamma}{1-\gamma} \mathbb{E}_{s_{t+1} \sim \mathcal{P}} \left[ C \sqrt{2D_{KL}(\pi_{\text{old}}(\cdot | s_{t+1}) | \pi_{\text{new}}(\cdot | s_{t+1}))} \right] \quad (6)$$

where  $\gamma$  is the discount factor,  $\mathcal{P}(s_t, \mathbf{u}_t, \cdot)$  is the transition function,  $C = \max_{\mathbf{u}_{t+1}} \Delta_Q(s_{t+1}, \mathbf{u}_{t+1})$ , and  $\Delta_Q(s_{t+1}, \mathbf{u}_{t+1}) \triangleq |Q_{tot}^{\text{old}}(s_{t+1}, \mathbf{u}_{t+1}) - \sum_i \text{OPT}_i(Q_i^{\text{old}}(o_{t+1}^i, u_{t+1}^i), s_{t+1})|$ .

The proof is provided in Appendix B, where we also analyze how this error affects the general convergence properties of policy

**Algorithm 1** ME-IGM

```

1: Input:  $\theta, \phi$ 
2: Initiate:  $\theta^- \leftarrow \theta, D \leftarrow \emptyset$ 
3: for each iteration do
4:   for each environment step do
5:     get  $\pi_i(u_t^i | o_t^i, s_t)$  via Eq. 10
6:      $u_t^i \sim \pi_i(u_t^i | o_t^i, s_t), \forall i \in \{1, \dots, n\}$ 
7:      $s_{t+1} \sim \mathcal{P}(s_t, \mathbf{u}_t, \cdot)$ 
8:      $D \leftarrow D \cup \{(s_t, \mathbf{u}_t, r(s_t, \mathbf{u}_t), s_{t+1})\}$ 
9:   end for
10:  for each gradient step do
11:    calculate  $\mathcal{T}_\lambda^{\pi_{jt}} Q_{tot}$  via Eq.
12:     $\theta \leftarrow \theta - \lambda_\theta \nabla_\theta L(\theta)$ 
13:     $\phi \leftarrow \phi - \lambda_\phi \nabla_\phi L(\phi)$ 
14:     $\omega \leftarrow \omega - \lambda_\omega \nabla_\omega L(\omega)$ 
15:     $\theta^- \leftarrow \tau\theta + (1 - \tau)\theta^-$ 
16:  end for
17: end for

```

**Table 1:** (a) presents the payoff matrix for a one-step matrix game involving two agents: the row player and the column player. Each agent has three possible actions: {A, B, C}. The rewards for each joint action are given in a  $3 \times 3$  matrix. This payoff structure is non-monotonic, as the optimal action for one player depends on the action chosen by its teammate. (b)–(d) illustrate the results of ME-QMIX, FOP, and QMIX, respectively. The first row and first column display the policies adopted by the row and column players after  $10k$  training steps. The  $3 \times 3$  matrix in the bottom right corner represents the global Q-function  $Q_{tot}$  learned by each algorithm. We can see that ME-QMIX is the only method that assigns the highest probability to selecting the optimal joint action, highlighting its effectiveness in overcoming the misalignment issue and achieving optimal coordination.  $\epsilon$  in (d) denotes the exploration rate used in  $\epsilon$ -greedy exploration.

| (a) Payoff Matrix |          |     |     | (b) Result of ME-QMIX (Ours) |            |               |            |
|-------------------|----------|-----|-----|------------------------------|------------|---------------|------------|
|                   | A        | B   | C   |                              | 1.         | 0.            | 0.         |
| A                 | <b>8</b> | -12 | -12 | 1.                           | <b>8</b>   | -12           | -14        |
| B                 | -12      | 0   | 0   | 0.                           | -8         | -28           | -30        |
| C                 | -12      | 0   | 0   | 0.                           | -10        | -30           | -32        |
| (c) Result of FOP |          |     |     | (d) Result of QMIX           |            |               |            |
|                   | .21      | .39 | .40 |                              | $\epsilon$ | $1-2\epsilon$ | $\epsilon$ |
| .21               | <b>8</b> | -12 | -12 | $\epsilon$                   | -8.5       | -8.5          | -8.5       |
| .40               | -12      | 0   | 0   | $1-2\epsilon$                | -8.5       | <b>0.2</b>    | 0.2        |
| .39               | -12      | 0   | 0   | $\epsilon$                   | -8.5       | 0.2           | 0.2        |

improvement. This motivates us to minimize the factor  $C$  and adopt the following objective to train the OPTs.

$$\min_{\text{OPT}_{1:n}} \mathbb{E}_{s, \mathbf{u} \sim \rho_{\pi_{jt}^{new}}} [\Delta_Q(s, \mathbf{u})^2], \quad (7)$$

where  $\pi_{jt}^{new} = (\pi_1^{new}, \dots, \pi_n^{new})$  and  $\pi_i^{new}$  is defined in Eq. (5).

**5 THE OVERALL FRAMEWORK: ME-IGM**

Figure 3 illustrates the framework of ME-IGM, which consists of three types of networks. Agents 1 to  $n$  share a network  $f_\theta$  for predicting local Q-values. The Mixer network, which aggregates local Q-values and manages the credit assignment among agents, is parameterized through its associated hyper-network as  $\text{Mixer}_\theta$ <sup>4</sup>.  $[\text{OPT}_{i,\phi}]_{i=1}^n$  are  $n$  order-preserving transformations, and  $\text{OPT}_{i,\phi}(f_\theta(o_t^i), s_t)$  defines a local stochastic policy for agent  $i$ . The update rules for these networks are described as follows.

ME-IGM optimizes two loss functions to update  $\theta$  and  $\phi$  separately.  $\text{Mixer}_\theta$  and  $f_\theta$  are trained by minimizing  $L(\theta)$ :

$$\mathbb{E}_{\tau_t} \left[ \frac{1}{2} \left( Q_{tot\theta}(s_t, \mathbf{u}_t) - \widehat{\mathcal{T}}^{\pi_{jt}} Q_{tot}(s_t, \mathbf{u}_t) \right)^2 \right] \quad (8)$$

where the variables are defined as:

$$\tau_t = \{o_{t:t+1}^{1:n}, u_{t:t+1}^{1:n}, s_{t:t+1}, r_t\}$$

$$Q_{tot\theta}(\mathbf{u}, s) = \text{Mixer}_\theta(f_\theta(o^1, u^1), \dots, f_\theta(o^n, u^n), s)$$

and the estimated Q-target is:

$$\widehat{\mathcal{T}}^{\pi_{jt}} Q_{tot}(s_t, \mathbf{u}_t) = r_t + \gamma Q_{tot, \theta^-}(s_{t+1}, \mathbf{u}_{t+1})$$

As a common practice, we utilize a target Q network  $Q_{tot, \theta^-}$  to stabilize training, with  $\theta^-$  being the exponentially weighted moving average of  $\theta$ . Notably, local actions  $u_{t:t+1}^{1:n}$  are sampled from local stochastic policies  $\pi_i$  (as defined in Equation (10)) during roll-out, rather than selecting the argmax of local Q-functions. This offers two advantages: (1) the softmax policy function enhances exploration, and (2) incorporating state  $s$  in the sampling process allows for better utilization of global information during centralized training.

Further, to train  $[\text{OPT}_{i,\phi}]_{i=1}^n$ , we adopt a sample-based version of Equation (7), denoted as  $L(\phi)$ :

$$\mathbb{E}_{\tau_t} \left[ \left( \sum_{i=1}^n \text{OPT}_{i,\phi}(f_\theta(o_t^i, u_t^i), s_t) - Q_{tot,\theta}(s_t, \mathbf{u}_t) \right)^2 \right] \quad (9)$$

OPTs define the local policies for each agent as follows:

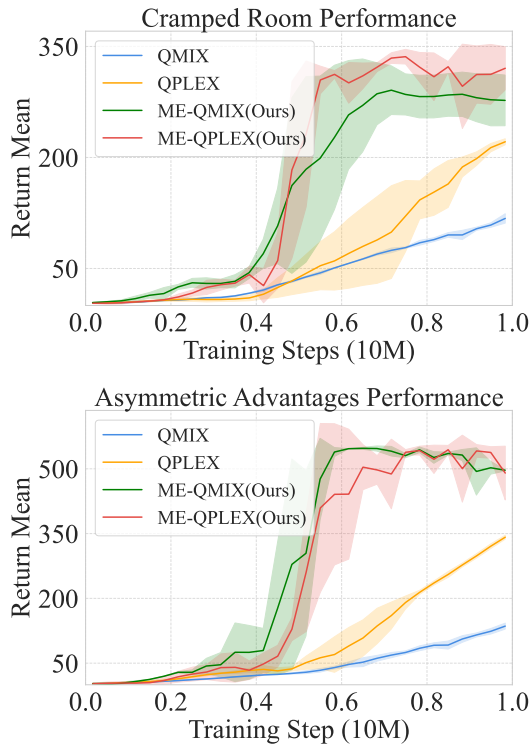
$$\pi_i(u^i | o^i, s) = \text{softmax}(\text{OPT}_{i,\phi}(f_\theta(o^i), s) / \alpha_\omega) \quad (10)$$

$\alpha_\omega$  is the temperature and, similar to Harnoja et al. [9], can be updated by minimizing the following loss function:

$$L(\omega) = \mathbb{E}_{\tau_t} [-\alpha_\omega \log \pi_{jt}(\mathbf{u}_t | s_t) - \alpha_\omega \mathcal{H}], \quad (11)$$

where  $\mathcal{H}$  is a predefined target entropy. Note that the local policies (i.e.,  $\pi_i$ ) are only used to collect training roll-outs  $\tau_t$ . During execution, we would select the argmax of  $f_\theta$ , which does not require any centralized information (e.g.,  $s$ ), since OPTs preserve the order of the input vectors. The pseudo code of ME-IGM algorithm is presented in Algorithm 1.

<sup>4</sup>Our algorithm is compatible with any Mixer network design that satisfies the IGM condition. For instance, during evaluation, we utilize the Mixer networks of QMIX and QPLEX, leading to two variants of our algorithm: ME-QMIX and ME-QPLEX.



**Figure 4: The mean return and standard deviation of QMIX, QPLEX, ME-QMIX, ME-QPLEX in Overcooked. It shows that ME-QMIX and ME-QPLEX achieve higher returns and exhibit faster convergence, compared to QMIX and QPLEX which do not adopt maximum entropy.**

## 6 EXPERIMENTS

In this section, we first evaluate our method on a classic matrix game, demonstrating its capability to learn a globally optimal policy despite the non-monotonic nature of the overall payoff matrix. Next, we compare our algorithm against a set of baselines on the SMAC-v2 and Overcooked benchmark, highlighting its state-of-the-art performance in cooperative MARL. Moreover, we conduct ablation studies to evaluate the benefits of entropy-driven exploration, assess the necessity of incorporating the order-preserving transformation, and analyze the impact of key hyperparameters.

### 6.1 Evaluation on Matrix Games

We first demonstrate the effectiveness of our algorithm in a classic one-step matrix game, shown as Table 1. It is a non-monotonic payoff matrix commonly used in previous studies [24, 26]. We evaluate ME-QMIX, a variant of ME-IGM built upon QMIX, by comparing its performance against two state-of-the-art MARL methods: QMIX [22], a value decomposition MARL method, and FOP [28], a Maximum Entropy MARL method. We are concerned with whether the joint action with the highest selection probability corresponds to the globally maximal Q value, ensuring optimal returns when executing a deterministic policy.

As shown in Table 1, QMIX suffers from high estimation error in payoff values of this non-monotonic game, due to its monotonicity constraints (as introduced in Section 2.3). Its deterministic policy converges to a suboptimal joint action (B, B) with a reward of 0. FOP, which employs a network architecture similar to QPLEX, is capable of accurately estimating value functions (i.e., the payoff). However, its policy tends to select suboptimal joint actions more frequently due to the misalignment between local policies and the maximum global Q-values, as discussed in Section 3. This observation highlights that the misalignment between learned policies and optimal joint actions is not confined to IGM-based approaches. ME-QMIX exhibits smaller estimation errors for the payoff of optimal joint actions (which is our primary concern) while having larger errors for suboptimal ones. Crucially, the joint action with the maximum Q-value in ME-QMIX aligns with the true optimal joint action, and the local policy assigns the highest probability to selecting this optimal joint action, avoiding the misalignment issue through the use of OPTs.

### 6.2 Evaluation on Overcooked

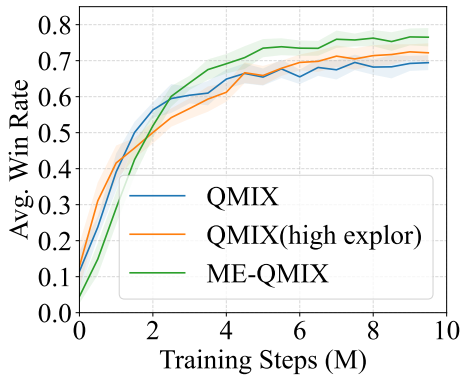
To demonstrate the advantage of employing the principle of maximum entropy, we compare ME-IGM with two leading value decomposition MARL methods (which do not involve maximum entropy): QMIX and QPLEX, on the Overcooked benchmark [3] – a widely used coordination-focused multi-agent environment. In particular, we evaluate them in two representative layouts: cramped room and asymmetric advantages. These scenarios pose distinct challenges, including close collaboration under spatial constraints and asymmetric advantages in task division, and require coordinated exploration. For each scenario, we conduct repeated experiments with different random seeds to get the training curves of 10 million environment steps. As shown in Figure 4, ME-QMIX and ME-QPLEX exhibit significantly faster convergence and achieve higher performance levels in the early stages of training compared to their counterparts, QMIX and QPLEX.

### 6.3 Evaluation on SMAC-v2

We further conduct a comprehensive evaluation of ME-IGM on the SMACv2 [5] benchmark across 15 scenarios, comparing its performance against 8 baselines. The results, as shown in Table 2, are obtained by training each algorithm on each scenario for 10 millions of environment steps. Each value in the table represents the average performance over three runs with different random seeds. ME-IGM outperforms the baseline algorithms across all scenarios by a significant margin. MAPPO and IPPO, as on-policy algorithms, have lower sampling efficiency, resulting in lower win rates within 10M environment steps. QMIX and QPLEX, lacking exploration, tend to converge to local optimum. HASAC [16] and FOP perform poorly across all scenarios. Surprisingly, ME-IGM achieves performance on par with (or even surpassing) IL+MARL methods, despite not relying on expert opponent demonstrations or opponent modeling. The results of MAPPO, IPPO, QMIX, QPLEX, IMAX-PPO, InQ are taken from [2]. A detailed comparison with actor-critic baselines is presented in the Appendix.

**Table 2: Average win rates of different algorithms on SMAC-v2. The results we recorded demonstrate that our method, ME-IGM, outperforms all previous baseline algorithms and achieves performance comparable to IL+MARL methods. Note that the results for MAPPO, IPPO, QMIX, QPLEX, IMAX-PPO, and InQ are cited directly from Bui et al. [2] as mean values; therefore, no standard deviations are reported for these baselines.**

| Task    | Scenario | MARL        |             |             |       |       |      | ME-IGM (Ours)   |                 | IL+MARL     |      |
|---------|----------|-------------|-------------|-------------|-------|-------|------|-----------------|-----------------|-------------|------|
|         |          | MAPPO       | IPPO        | QMIX        | QPLEX | HASAC | FOP  | ME-QMIX         | ME-QPLEX        | IMAX-PPO    | InQ  |
| Protoss | 5_vs_5   | 58.0        | 54.6        | <u>70.2</u> | 53.3  | 22.9  | 37.8 | <b>75.9±3.0</b> | 69.8±1.8        | 68.1        | 78.7 |
|         | 10_vs_10 | 58.3        | 58.0        | 69.0        | 53.7  | 12.1  | 8.6  | <u>78.5±2.7</u> | <b>78.7±3.4</b> | 59.6        | 79.8 |
|         | 10_vs_11 | 18.2        | 20.3        | 42.5        | 22.8  | 5.7   | 0.4  | <b>50.4±2.9</b> | <u>44.9±3.2</u> | 21.3        | 48.7 |
|         | 20_vs_20 | 38.1        | 44.5        | <u>69.7</u> | 27.2  | -     | 2.0  | <b>73.9±3.0</b> | 38.3±4.0        | 76.3        | 80.6 |
|         | 20_vs_23 | 5.1         | 4.1         | 16.5        | 4.8   | -     | 0.3  | <b>23.4±1.0</b> | <u>18.7±1.3</u> | 11.8        | 24.2 |
| Terran  | 5_vs_5   | 52.0        | 56.2        | 58.4        | 70.0  | 37.9  | 57.1 | <b>70.2±2.6</b> | <u>70.0±4.8</u> | 53.3        | 69.9 |
|         | 10_vs_10 | 58.1        | 57.3        | 65.8        | 66.1  | 20.4  | 15.8 | <b>72.6±1.8</b> | <u>69.4±0.2</u> | 58.4        | 72.2 |
|         | 10_vs_11 | 28.6        | 31.0        | 39.4        | 41.4  | 9.1   | 7.5  | <u>49.6±1.6</u> | <b>51.0±0.4</b> | 28.4        | 53.9 |
|         | 20_vs_20 | 52.8        | 49.6        | <u>57.6</u> | 23.9  | -     | 0.2  | <b>59.3±1.3</b> | 45.1±7.2        | 35.9        | 65.4 |
|         | 20_vs_23 | 11.2        | 10.0        | 10.0        | 7.0   | -     | 0.0  | <b>18.7±0.7</b> | 17.4±2.9        | 4.7         | 17.7 |
| Zerg    | 5_vs_5   | 41.0        | 37.2        | 37.2        | 47.8  | 29.1  | 35.2 | <u>50.8±4.6</u> | <b>52±3.9</b>   | 48.6        | 55.0 |
|         | 10_vs_10 | 39.1        | <b>49.4</b> | 40.8        | 41.6  | 17.9  | 8.0  | <u>49.1±0.6</u> | 42.9±1.8        | 50.6        | 57.6 |
|         | 10_vs_11 | 31.2        | 26.0        | 28.0        | 31.1  | 14.0  | 1.5  | 32.6±1.7        | <b>45.3±0.6</b> | <u>34.8</u> | 41.5 |
|         | 20_vs_20 | 31.9        | 31.2        | 30.4        | 15.8  | -     | 0.2  | <u>34.3±5.2</u> | <b>42.1±5.1</b> | 26.7        | 43.3 |
|         | 20_vs_23 | <u>15.8</u> | 8.3         | 10.1        | 6.7   | -     | 0.2  | 14.8±6.4        | <b>17.5±3.2</b> | 8.2         | 21.3 |



**Figure 5: Ablation study on exploration strategies in QMIX. Comparison between ME-QMIX and a modified QMIX with extended epsilon annealing on the Protoss 5v5 map. The results show that simply increasing epsilon-greedy exploration is insufficient for improving performance. In contrast, ME-QMIX enables QMIX agents to perform more structured exploration through the maximum entropy framework, leading to higher returns.**

### 6.4 Ablation Study

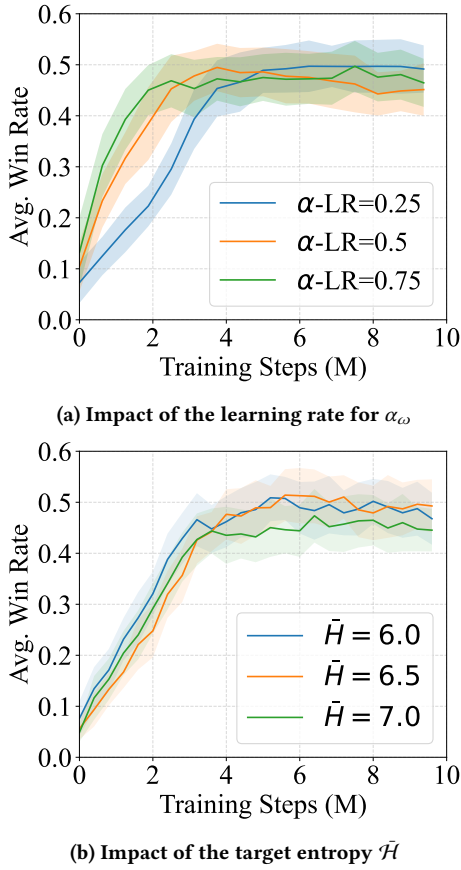
**Benefits of entropy-driven exploration:** To further demonstrate the benefits of using the maximum entropy framework for exploration, we compare ME-QMIX with a modified version of QMIX, which incorporates significantly higher exploration by applying an extended epsilon annealing period in the epsilon-greedy method,

**Table 3: Our algorithm (+OPT) vs. its two ablated versions: +Entropy, obtained by removing OPTs, and QMIX, obtained by further removing softmax local policies. Results on three SMAC-v2 scenarios show that removing OPTs (+Entropy) leads to lower win rates due to misalignment between local and global objectives, while incorporating OPTs ensures consistent and monotonic policy improvement.**

|                     | QMIX        | +Entropy    | +OPT        |
|---------------------|-------------|-------------|-------------|
| <b>Protoss 5vs5</b> | 0.68 ± 0.04 | 0.73 ± 0.05 | 0.74 ± 0.04 |
| <b>Terran 5vs5</b>  | 0.68 ± 0.05 | 0.67 ± 0.04 | 0.70 ± 0.05 |
| <b>Zerg 5vs5</b>    | 0.41 ± 0.05 | 0.40 ± 0.04 | 0.48 ± 0.05 |

on Protoss 5v5. Figure 5 illustrates that simply increasing the epsilon annealing time for exploration does not improve QMIX’s performance. ME-QMIX achieves superior performance, with the systematic exploration enabled by the maximum entropy framework.

**The necessity of using OPTs:** In Table 3, we compare QMIX with its two variants: +Entropy and +OPT. Both variants learn local stochastic policies based on Q-functions trained with the QMIX objective. Specifically, +OPT corresponds to our proposed ME-IGM, while +Entropy removes the OPTs from ME-IGM (i.e., no OPTs in Equations (4), (9), and (10)). We test each algorithm on three SMAC-v2 scenarios across four random seeds, reporting the average win rate along with its standard deviation in the table. Without OPTs, +Entropy suffers from the misalignment problem discussed in Section 3 and fails to ensure monotonic policy improvement, as analyzed in Section 4.2, leading to a performance decline in two out



**Figure 6: The impact of different hyperparameters for training  $\alpha_\omega$  on the performance of ME-QMIX. (a) A lower learning rate for  $\alpha_\omega$  corresponds to increased exploration. (b) ME-QMIX’s performance is not sensitive to the choice of the target entropy  $\tilde{H}$  in Equation (10).**

**Table 4: Ablation study on the effect of the Order-Preserving Transformation (OPT). We compare ME-QMIX (OPT) with a variant that replaces OPT layers with a standard 3-layer MLP, which does not preserve the order of input values. Results on three SMAC-v2 scenarios show that ME-QMIX (OPT) consistently achieves higher win rates, highlighting that enforcing order preservation is essential for mitigating local-global misalignment and achieving stable policy improvement.**

|                    | ME-QMIX (MLP)   | ME-QMIX (OPT)   |
|--------------------|-----------------|-----------------|
| <b>Protoss 5v5</b> | 0.58 $\pm$ 0.03 | 0.74 $\pm$ 0.04 |
| <b>Terran 5v5</b>  | 0.61 $\pm$ 0.04 | 0.70 $\pm$ 0.05 |
| <b>Zerg 5v5</b>    | 0.41 $\pm$ 0.04 | 0.48 $\pm$ 0.05 |

of three scenarios. On the other hand, +OPT achieves a substantial performance improvement across all evaluated scenarios.

To further validate the design of the order-preserving transformation, we replace the OPT layers in ME-QMIX with a 3-layer

MLP, which does not enforce non-negative weights, referring to this variant as ME-QMIX (MLP). As shown in Table 4, our method, ME-QMIX (OPT), consistently outperforms ME-QMIX (MLP) across three different SMAC-v2 scenarios. This underscores the importance of the order-preserving transformation (i.e., Eq. (2)) in resolving the misalignment between local policies and the maximum global Q-value, thereby enhancing overall performance.

**The impact of hyperparameters:** Compared to QMIX, ME-QMIX introduces two additional hyperparameters: the learning rate and the target entropy  $\tilde{H}$  for updating  $\alpha_\omega$  in Eq. (10). We set the other hyperparameters of ME-QMIX following the design choices of QMIX. To analyze the impact of these additional parameters, we conduct a parameter scan on the learning rate of  $\alpha_\omega$  and the target entropy  $\tilde{H}$ . Figure 6a illustrates that the learning rate of  $\alpha_\omega$  influences the exploration level. A lower learning rate results in higher policy entropy at the beginning of training, promoting exploration. While this initially leads to weaker performance (e.g., the blue curve in Figure 6a), it eventually achieves higher returns. Furthermore, the results in Figure 6b suggest that the algorithm’s performance is not sensitive to the choice of the target entropy. Notably, in this study, we did not fine-tune configurations for each individual testing scenario. Please check the Appendix for the tables of hyperparameters.

## 7 CONCLUSION

In this paper, we introduce ME-IGM, a novel maximum entropy MARL approach that applies the CTDE framework and satisfies the IGM condition. We identify and address the misalignment problem between local policies and the maximum global Q-value (in maximum entropy MARL) by introducing an order-preserving transformation (OPT). We further propose a theoretically sound and straightforward objective for updating these transformation operators. Empirically, we demonstrate the effectiveness of ME-IGM on matrix games and its state-of-the-art performance on the challenging Overcooked and SMAC-v2 benchmarks. However, as a value-based method, ME-IGM is currently limited to tasks with discrete action spaces. In future work, we plan to extend its use to tasks involving continuous action spaces.

## REFERENCES

- [1] Zafarali Ahmed, Nicolas Le Roux, Mohammad Norouzi, and Dale Schuurmans. 2019. Understanding the impact of entropy on policy optimization. In *International conference on machine learning*. PMLR, 151–160.
- [2] The Viet Bui, Tien Anh Mai, and Thanh Hong Nguyen. 2024. Mimicking To Dominate: Imitation Learning Strategies for Success in Multiagent Games. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=06JRFVK880>
- [3] Micah Carroll, Rohin Shah, Mark K Ho, Tom Griffiths, Sanjit Seshia, Pieter Abbeel, and Anca Dragan. 2019. On the utility of learning about humans for human-ai coordination. *Advances in neural information processing systems* 32 (2019).
- [4] Djork-Arné Clevert. 2015. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289* (2015).
- [5] Benjamin Ellis, Jonathan Cook, Skander Moalla, Mikayel Samvelyan, Mingfei Sun, Anuj Mahajan, Jakob N Foerster, and Shimon Whiteson. 2022. Smacv2: An improved benchmark for cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2212.07489* (2022).
- [6] Benjamin Eysenbach and Sergey Levine. 2021. Maximum entropy rl (provably) solves some robust rl problems. *arXiv preprint arXiv:2103.06257* (2021).
- [7] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [8] Fenggen Guo and Zizhao Wu. 2022. Learning maximum entropy policies with QMIX in cooperative MARL. In *2022 IEEE 2nd International Conference on Electronic Technology, Communication and Information (ICETCI)*. IEEE, 357–361.

- [9] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [10] Zichen He, Lu Dong, Chunwei Song, and Changyin Sun. 2022. Multiagent soft actor-critic based hybrid motion planner for mobile robots. *IEEE transactions on neural networks and learning systems* (2022).
- [11] Jian Hu, Siyang Jiang, Seth Austin Harding, Haibin Wu, and Shih-wei Liao. 2021. Rethinking the implementation tricks and monotonicity constraint in cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2102.03479* (2021).
- [12] Taehyeon Kim, Jaehoon Oh, NakYil Kim, Sangwook Cho, and Se-Young Yun. 2021. Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation. *arXiv preprint arXiv:2105.08919* (2021).
- [13] Jakub Grudzien Kuba, Ruiqing Chen, Muning Wen, Ying Wen, Fanglei Sun, Jun Wang, and Yaodong Yang. 2021. Trust region policy optimisation in multi-agent reinforcement learning. *arXiv preprint arXiv:2109.11251* (2021).
- [14] Sergey Levine. 2018. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909* (2018).
- [15] Fanqi Lin, Shiyu Huang, Tim Pearce, Wenzhe Chen, and Wei-Wei Tu. 2023. Tizer0: Mastering multi-agent football with curriculum learning and self-play. *arXiv preprint arXiv:2302.07515* (2023).
- [16] Jiarong Liu, Yifan Zhong, Siyi Hu, Haobo Fu, Qiang Fu, Xiaojun Chang, and Yaodong Yang. 2023. Maximum Entropy Heterogeneous-Agent Mirror Learning. *arXiv preprint arXiv:2306.10715* (2023).
- [17] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* 30 (2017).
- [18] Frans A Oliehoek and Christopher Amato. 2016. *A concise introduction to decentralized POMDPs*. Springer.
- [19] Bei Peng, Tabish Rashid, Christian Schroeder de Witt, Pierre-Alexandre Kamieny, Philip Torr, Wendelin Böhrer, and Shimon Whiteson. 2021. Facmac: Factored multi-agent centralised policy gradients. *Advances in Neural Information Processing Systems* 34 (2021), 12208–12221.
- [20] Mark S Pinsker. 1964. Information and information stability of random variables and processes. *Holden-Day* (1964).
- [21] Yuan Pu, Shaochen Wang, Rui Yang, Xin Yao, and Bin Li. 2021. Decomposed soft actor-critic method for cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2104.06655* (2021).
- [22] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *The Journal of Machine Learning Research* 21, 1 (2020), 7234–7284.
- [23] Elena Smirnova and Elvis Dohmatob. 2019. On the Convergence of Approximate and Regularized Policy Iteration Schemes. *arXiv preprint arXiv:1909.09621* (2019).
- [24] Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International conference on machine learning*. PMLR, 5887–5896.
- [25] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 2017. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296* (2017).
- [26] Jianhao Wang, Zhizhou Ren, Terry Liu, Yang Yu, and Chongjie Zhang. 2020. Qplex: Duplex dueling multi-agent q-learning. *arXiv preprint arXiv:2008.01062* (2020).
- [27] Chao Yu, Akash Velu, Eugene Vinytsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems* 35 (2022), 24611–24624.
- [28] Tianhao Zhang, Yueheng Li, Chen Wang, Guangming Xie, and Zongqing Lu. 2021. Fop: Factorizing optimal joint policy of maximum-entropy multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 12491–12500.
- [29] Yifan Zhong, Jakub Grudzien Kuba, Siyi Hu, Jiaming Ji, and Yaodong Yang. 2023. Heterogeneous-Agent Reinforcement Learning. *arXiv preprint arXiv:2304.09870* (2023).
- [30] Hanhan Zhou, Tian Lan, and Vaneet Aggarwal. 2022. Pac: Assisted value factorization with counterfactual predictions in multi-agent reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 15757–15769.