

Think How Your Teammates Think: Active Inference Can Benefit Decentralized Execution

Hao Wu^{1, 2*}, Shoucheng Song^{1, 2*}, Chang Yao^{1, 2}, Sheng Han^{1, 2},
Huaiyu Wan^{1, 2}, Youfang Lin^{1, 2}, Kai Lv^{1, 2†}

¹School of Computer Science & Technology, Beijing Jiaotong University, Beijing, China

²Beijing Key Laboratory of Traffic Data Mining and Embodied Intelligence, Beijing, China
{wuhao-, insis_songsc, yaochang, shhan, hywan, yflin, lvkai}@bjtu.edu.cn

Abstract

In multi-agent systems, explicit cognition of teammates’ decision logic serves as a critical factor in facilitating coordination. Communication (i.e., “*Tell*”) can assist in the cognitive development process by information dissemination, yet it is inevitably subject to real-world constraints such as noise, latency, and attacks. Therefore, building the understanding of teammates’ decisions without communication remains challenging. To address this, we propose a novel non-communication MARL framework that realizes the construction of cognition through local observation-based modeling (i.e., “*Think*”). Our framework enables agents to model teammates’ **active inference** process. At first, the proposed method produces three teammate portraits: perception-belief-action. Specifically, we model the teammate’s decision process as follows: 1) Perception: observing environments; 2) Belief: forming beliefs; 3) Action: making decisions. Then, we selectively integrate the belief portrait into the decision process based on the accuracy and relevance of the perception portrait. This enables the selection of cooperative teammates and facilitates effective collaboration. Extensive experiments on the SMAC, SMACv2, MPE, and GRF benchmarks demonstrate the superior performance of our method.

Introduction

Multi-agent reinforcement learning (MARL) has garnered significant attention due to its wide applications in fields such as autonomous driving (Kiran et al. 2021), smart grids (Roesch et al. 2020), and transportation (Lee, Chung, and Sohn 2019). In decentralized systems, the lack of cognition regarding teammates’ decision logic may induce miscoordination among agents and result in suboptimal policies.

To address the mentioned issue, one intuitive approach is “*Tell agent how its teammates act*”, which can be implemented through communication mechanisms. The communication (i.e., “*Tell*”) methods can facilitate the understanding of teammates’ behaviors by exchanging decision-relevant messages (Wang et al. 2019; Yuan et al. 2022; Sun et al. 2023a, 2024). However, communication effectiveness may be limited under some conditions, including limited

*These authors contributed equally.

†Corresponding Author: Kai Lv (lvkai@bjtu.edu.cn).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

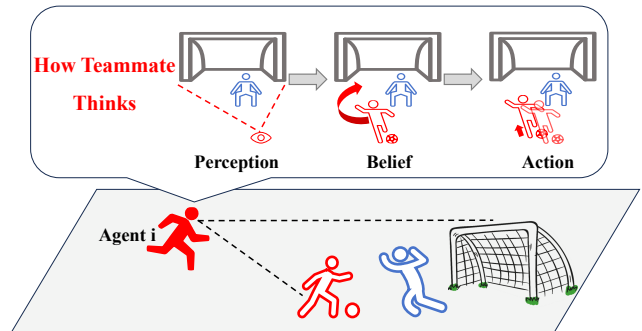


Figure 1: Modeling the Active Inference Process of the Teammate. In this scenario, agent i models the perception-belief-action involved in its teammate’s active inference process when facing the goalkeeper. This allows i to obtain its teammate’s decision-relevant information and achieve effective collaboration.

bandwidth, high latency, and significant noise (Tung et al. 2021; Hu et al. 2023; Song et al. 2025). Therefore, we explore the method for developing the cognition of teammates’ decisions in the absence of communication.

In contrast to directly telling an agent how its teammates act, we advocate that agents could engage in “*Think how your teammates think*”. Specifically, the agent actively builds comprehension of teammates’ decisions. To achieve this, the direct method is to model teammates’ decision process. However, existing agent modeling methods fail to fulfill this. On the one hand, some methods rely on access to other agents’ trajectories during the modeling (Rabinowitz et al. 2018; Zintgraf et al. 2021), which is unavailable during decentralized execution. On the other hand, some methods typically only enable a single agent to model other agents that possess fixed parameters (Xie et al. 2021; Papoudakis, Christianos, and Albrecht 2021; Yu, Jiang, and Lu 2024). This configuration imposes an upper bound on the system’s collaborative efficiency, preventing the team from learning more optimal policies. Furthermore, these methods only model incomplete decision components (e.g., behaviors or intentions), risking inaccuracies from discrepancies

between the model and the actual situation. To address this, we rely solely on local observations to model teammates’ complete decision processes, reducing modeling inaccuracies during decentralized execution.

Inspired by human brain decision-making mechanisms and active inference theory (Friston et al. 2016), we model the teammate’s decision process as an active inference process comprising perception-belief-action. In this process, the teammate perceives environments, forms beliefs, and then takes actions by integrating perceptions and beliefs. Consequently, we employ local observation-based modeling (i.e., “*Think*”) to acquire teammates’ perception-belief-action (i.e., “*How your teammates think*”). Figure 1 provides an illustrative example of this modeling process.

In this paper, we propose a novel non-communication framework for modeling the Active Inference of teammates in MARL (AIM). The framework consists of two parts: At first, we develop a modeling method to model teammates’ three portraits: perception-belief-action, solely based on local observations. Meanwhile, the perception and belief portraits are optimized by minimizing the discrepancy between predicted and actual actions. Then, we propose a dual-filter mechanism to enhance teammates’ cognition utilization. This mechanism features selective collaboration by choosing teammates whose modeled portraits have high accuracy. Additionally, by considering the perception relevance among agents, we adopt an attention module to dynamically integrate teammates’ belief portraits, thereby optimizing the decision process. Our proposed method demonstrates significant improvement in tasks within SMAC (Samvelyan et al. 2019), SMACv2 (Ellis et al. 2024), MPE (Mordatch and Abbeel 2018) and GRF (Kurach et al. 2020).

Our contributions are outlined as follows:

- We replace “communication (i.e., *Tell*)” with “modeling (i.e., *Think*)”, enabling agents to construct the cognition of teammates’ decision logic without communication during decentralized execution.
- We propose an active inference framework to model teammates’ three portraits: perception-belief-action, to understand how they think.
- We introduce a dual filter that leverages the accuracy and relevance of perception portraits to select cooperative teammates.
- We conduct experiments on SMAC, SMACv2, MPE, and GRF. The results show that our method achieves optimal or near-optimal performance in most scenarios.

Related Works

Communication in MARL

Several communication methods, such as (Das et al. 2019; Ding, Huang, and Lu 2020; Yuan et al. 2022; Sun et al. 2023b; Sun 2024; Li et al. 2025; Yao et al. 2025), design communication networks that enable agents to exchange decision-relevant messages during the decentralized execution. However, in some real-world scenarios, limitations such as high noise, high latency, and low bandwidth often prevent these communication algorithms from performing

well (Tung et al. 2021; Hu et al. 2023; Song et al. 2025). Furthermore, the communication attack may introduce malicious information, disrupting agents’ decision-making and hindering collaboration (Xue et al. 2021; Zhu, Dastani, and Wang 2024). In comparison, our method adopts an alternative perspective. We propose a novel communication-free framework that models teammates’ active inference to comprehensively understand their decision logic.

Agent Modeling Methods

Methods based on agent modeling typically acquire agents’ behaviors, beliefs, or intentions. For example, (Rabinowitz et al. 2018; Yang et al. 2018; Tian et al. 2019; Zintgraf et al. 2021; Zhai et al. 2023) model agents’ psychological state and beliefs using the Theory of Mind and Bayesian reasoning. (He et al. 2016; Raileanu et al. 2018) infer agents’ policy based on the modeled agents’ observations and actions. (Papoudakis and Albrecht 2020) model agents’ policy as a latent distribution. However, these methods rely on a strong assumption that agents can directly access the modeled agent’ trajectories, which is infeasible during the decentralized execution.

Although (Papoudakis, Christianos, and Albrecht 2021; Xie et al. 2021; Yu, Jiang, and Lu 2024) model other agents solely based on local observations, they maintain teammates’ policies as fixed parameters. Nevertheless, collaboration performance becomes constrained by fixed teammate parameters, preventing team policy optimization. Moreover, the above methods only partially model other agents’ decision-making processes, introducing inevitable inaccuracies and negatively impacting cooperative performance owing to mismatches between the model and reality. In contrast, our framework models teammates’ complete decision-making process through three key aspects: perception-belief-action to increase modeling accuracy. In this way, our method enables agents to understand how they think.

Background

A fully cooperative multi-agent task can be modeled as a *Decentralized Partially Observable Markov Decision Process* (Dec-POMDP) (Oliehoek, Amato et al. 2016), represented as a tuple $\langle I, S, \{A_i\}_{i=1}^N, \{\Omega_i\}_{i=1}^N, O, \mathcal{T}, R, \gamma \rangle$. In this model, $I = \{1, \dots, N\}$ is the set of agents and N is the number of agents. S is the state space. For each agent $i \in I$, A_i is the individual action space, and $A = A_1 \times \dots \times A_N$ is the joint action space. Ω_i is the observation space for agent i . $O(o_i | s, i)$ is the observation function over local observations $o_i \in \Omega_i$ given state $s \in S$ and agent i . The state transition function $\mathcal{T}(s' | s, \mathbf{a})$ defines the probability of transitioning to state $s' \in S$ given current state $s \in S$ and joint action $\mathbf{a} \in A$. The reward function $R(s, \mathbf{a})$ gives the immediate reward, and $\gamma \in [0, 1)$ is the discount factor, used to balance long-term and short-term rewards. At each time step t , agent $i \in I$ receives a local observation $o_i \sim O(\cdot | s, i)$, selects an action $a_i \in A_i$, and the joint action $\mathbf{a} = \langle a_1, \dots, a_N \rangle$ results in a state transition and a reward $R(s, \mathbf{a})$. The goal of multi-agent system algorithms is to find a joint policy π to maximize the expected cumulative reward, formulated as: $V^\pi(s) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t R(s_t, \mathbf{a}_t)]$.

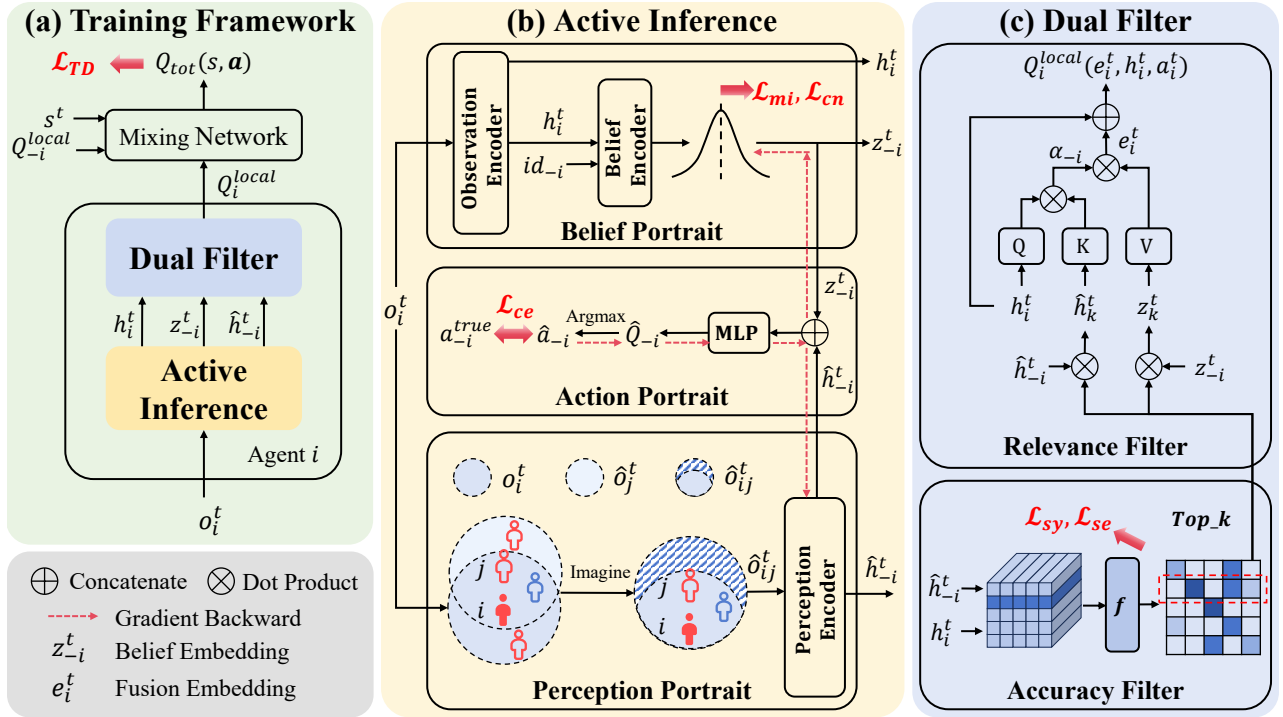


Figure 2: The overall framework of AIM. (a) The training framework comprises the agent network and the mixing network; (b) The active inference module, which includes perception portrait, belief portrait, and action portrait; (c) The dual filter module, consisting of the accuracy filter and the relevance filter.

Method

To build the cognition of teammates’ decision logic during decentralized execution, our core idea is to model their complete decision-making process. On the one hand, we introduce active inference to construct teammates’ three portraits: perception-belief-action, which represent “how teammates think”. On the other hand, due to modeling errors and the diversity of teammates, we devise a dual filter that dynamically integrates teammates’ belief portraits based on the accuracy and relevance of their perception portraits.

Teammate Portrait via Active Inference

In active inference, the teammate’s decision-making process is divided into three stages: perception of the environment, belief formation, and action execution. Perception serves as the foundational step, capturing the change of entities, which supports the following processes. Beliefs represent agents’ deeper understanding of the environment, serving as the basis for deriving action. The final decision, which is the action output, is based on both perception and belief and can be used as posterior information to optimize the perceptions and beliefs. Here, we give the details of the triple model process in AIM.

Perception Portrait To understand teammates’ behavior, agents should first understand what they have experienced. AIM aims to construct the world in teammates’ eyes, generating their local observations from a teammate-centered

perspective. In detail, for agent i , the perception \hat{o}_j^t of teammate j is constructed based on i ’s own observations o_i^t . The specific process involves setting each teammate’s position as the origin and recalculating the positions of the other agents relative to this origin, while also adding other relevant information. As shown in Figure 2, we only select the portion of o_i^t that intersects with \hat{o}_j^t as the perception portrait \hat{o}_{ij}^t . This process is essentially a viewpoint transformation operation. More details about the transformation can be found in Appendix B. We then use \hat{o}_{ij}^t as the input and apply a GRU to obtain the historical trajectory information \hat{h}_{ij}^t of the teammate j . Subsequently, we acquire \hat{h}_{-i}^t of all teammates.

Belief Portrait Belief serves as a higher-level abstraction of actions, serving as the core basis for policies. However, unlike the perception portrait, which is objective, the belief portrait typically exhibits high subjectivity and variability. In scenarios with limited observation, these characteristics become more pronounced, making it challenging to model beliefs from teammates’ perspectives accurately.

Therefore, we construct the belief portrait of teammates from the agent’s perspective. For agent i , we use its trajectory h_i^t and teammates’ index id_{-i} for modeling, as shown in Figure 2. The input (h_i^t, id_{-i}) is fed into the belief encoder to obtain the belief distribution $\mathcal{N}(\mu_i^t, \delta_i^t)$. Through reparameterization, we obtain the teammates’ belief portraits z_{-i}^t which should exhibit two characteristics: (1) Decision-support ability; (2) Stability over the short term.

To enhance the decision-support ability, we get inspired by (Yuan et al. 2022) and maximize the mutual information between teammates’ actions and belief portraits z_{-i}^t , conditioned on h_i^t and id_{-i} . Through variational inference, maximizing mutual information can be transformed into the following loss. The derivation can be found in Appendix A.

$$\mathcal{L}_{mi} = \mathbb{E} \left[\mathcal{D}_{KL} \left(p(z_{-i}^t | h_i^t, id_{-i}) \parallel q_{\xi}(z_{-i}^t | h_i^t, a_{-i}^t, id_{-i}) \right) \right], \quad (1)$$

where $\mathcal{D}_{KL}(\dots \parallel \dots)$ represents the Kullback-Leibler divergence, $q_{\xi}(z_{-i}^t | h_i^t, a_{-i}^t, id_{-i})$ is used as a variational distribution to approximate the conditional distribution $p(z_{-i}^t | h_i^t, id_{-i})$. Minimizing \mathcal{L}_{mi} is equivalent to maximizing the relevance between the belief portraits and selected actions.

To improve the stability of z_{-i}^t , we calculate the cosine similarity between the z_{-i}^t at adjacent time step, formalized as:

$$\mathcal{L}_{cn} = \mathbb{E} \left[- \frac{z_{-i}^{t-1} \cdot z_{-i}^t}{\|z_{-i}^{t-1}\| \|z_{-i}^t\|} \right]. \quad (2)$$

Action Portrait Action is the most direct outcome of active inference. The accuracy of the action serves as a critique of the modeling process, especially the coherence of the belief portrait. Moreover, joint actions induce environmental state transitions, which in turn influence teammates’ perception portraits. Hence, we optimize the perception and belief portrait by leveraging posterior action predictions. In AIM, we concatenate belief portraits z_{-i}^t and historical perception information \hat{h}_{-i}^t as the input of action prediction network, and obtain the imagined action distribution \hat{a}_{-i} . By minimizing the cross-entropy loss between \hat{a}_{-i} and the true actions a_{-i}^{true} , AIM optimizes the action modeling network and the perception-belief portrait.

$$\mathcal{L}_{ce} = - \sum_i a_{-i}^{true} \log \hat{a}_{-i}. \quad (3)$$

The complete loss for the teammate portrait via active inference is expressed as:

$$\mathcal{L}_{MD} = \lambda_{mi} \mathcal{L}_{mi} + \lambda_{cn} \mathcal{L}_{cn} + \lambda_{ce} \mathcal{L}_{ce}, \quad (4)$$

where λ_{mi} , λ_{cn} , and λ_{ce} represent the hyperparameters of the loss function, which are used to balance their effects.

Dual Filter of Accuracy and Relevance

Due to local observation, errors in the active inference process are unavoidable. Indiscriminately utilizing erroneous portraits of teammates can directly distort the agent’s comprehension of the current environment, resulting in non-cooperative behavior. Furthermore, since collaboration in multi-agent systems is often localized, it is redundant to incorporate the portraits of all teammates in decision-making. Hence, we propose a dual filter mechanism for selecting cooperative teammates, focusing on two aspects: the accuracy and the relevance of portraits.

Accuracy Filter Among triple portraits, the perception portrait is derived through perspective transformation, rendering it inherently partial. Therefore, an evaluation method is required to assess the accuracy of perception portraits.

Specifically, we learn a mapping $f : \mathbb{R}^h \mapsto \mathbb{R}$ that maps the perception portrait to an accuracy score. At time t , we simultaneously process the $N \times N$ portraits, constructing the evaluation matrix \mathcal{C}^t . N represents the number of agents. c_{ij}^t refers to the accuracy score of agent i ’s perception portrait \hat{h}_{ij}^t to agent j , formed as:

$$c_{ij}^t = \text{softmax}(f(\hat{h}_{ij}^t)), \quad (5)$$

where $f(\cdot)$ represents a two-layer MLP network. A higher value of c_{ij} signifies that the perception portrait of agent j by agent i is more accurate. The evaluation matrix \mathcal{C}^t should possess the following characteristics.

(1)**Mutual evaluation is similar.** For agents i and j , perception portraits of each other are essentially different perspectives on the same intersection of their observations.

(2)**Self-evaluation is highest.** The perception portrait of an agent’s own is guaranteed to be accurate, so the network’s evaluation of this result should be the highest.

(3)**High similarity gets a high score.** If an agent’s perception portrait is similar to the real observation of any other agent, it indicates high accuracy.

To satisfy the **Mutual evaluation is similar**, we introduce a symmetry loss \mathcal{L}_{sy} to optimize the evaluation matrix \mathcal{C} , which is formalized as:

$$\mathcal{L}_{sy} = \|\mathcal{C} - \mathcal{C}^T\|_{\mathcal{F}}, \quad (6)$$

where \mathcal{C}^T is the transpose of the matrix \mathcal{C} , and $\|\cdot\|_{\mathcal{F}}$ denotes the Frobenius-norm.

To enhance **Self-evaluation is highest**, we utilize a diagonal loss \mathcal{L}_{se} to maximize the accuracy score of true h_{ii}^t .

$$\mathcal{L}_{se} = - \sum_i c_{ii}. \quad (7)$$

Regarding **High similarity gets a high score**, deep neural networks exhibit the property that similar inputs produce similar outputs. Combined with characteristic (2), perception portraits similar to the true h_{ii}^t can be mapped to higher scores, thereby satisfying characteristic (3) without the additional loss function.

Having the evaluation matrix \mathcal{C} , we sample the *top-k* indices for each agent based on the accuracy to select potential teammates for the next filter.

Relevance Filter Building on the accuracy filter, we also need to apply the relevance filter from a decision-making perspective, selecting teammates more relevant to the agent for effective collaboration. Given that collaboration in multi-agent systems is often localized, we use the perception portrait as a proxy for relevance. As for how to utilize teammates’ portraits to aid decision-making, the most intuitive approach is to combine teammates’ action portraits to mitigate the lack of cognition about teammates. However, due to local observation, accurately modeling actions is challenging. Therefore, we instead combine belief portraits, leveraging higher-level behavioral bases to dilute the impact of single-step modeling errors.

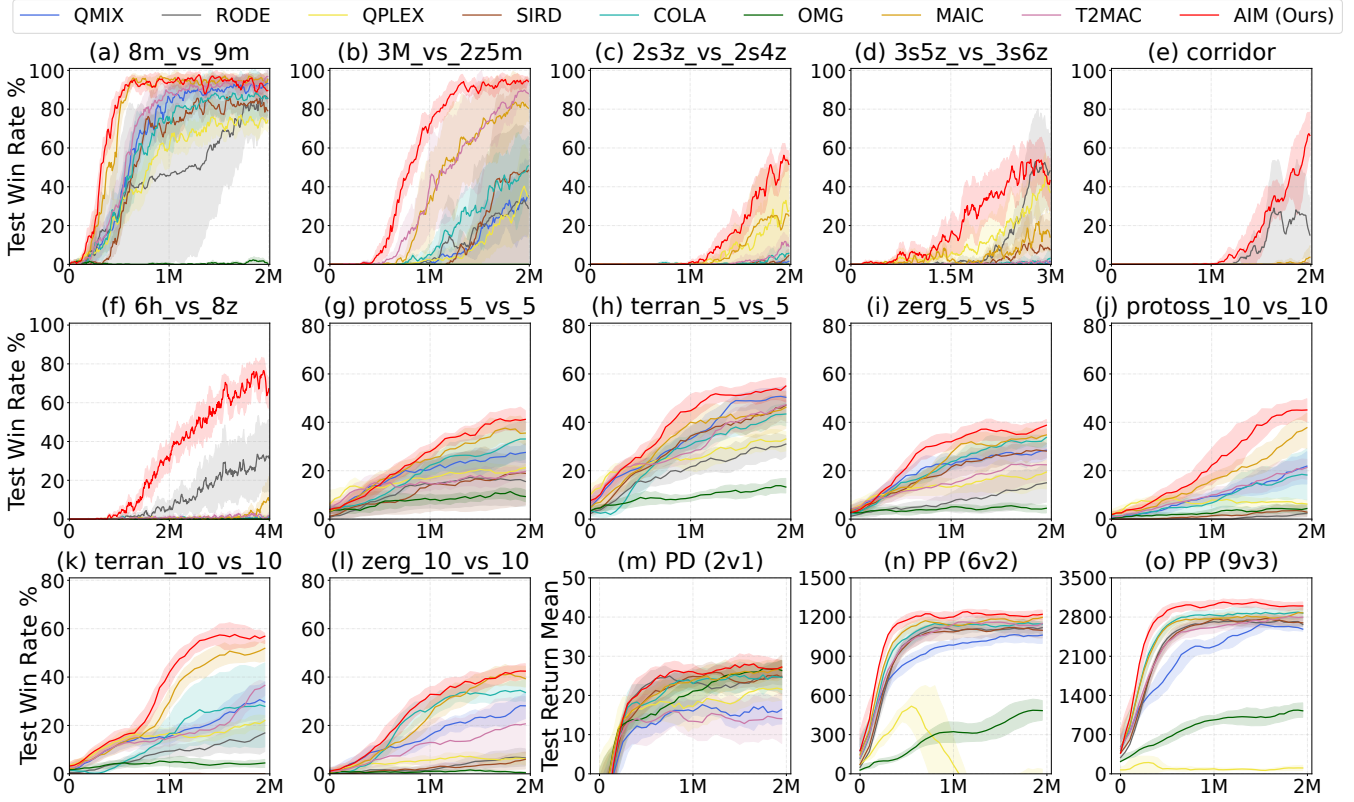


Figure 3: Performance comparison between AIM and baselines on SMAC, SMACv2, and MPE. (a)-(f) Six representative maps on SMAC. (g)-(l) Six tasks on SMACv2. (m)-(o) Three tasks on MPE.

Here, we apply the attention mechanism (Vaswani 2017) to achieve the relevance filter and fusion of belief portraits. In AIM, for agent i , we adopt the true perception history h_i^t as the query, the historical perception \hat{h}_k^t of the selected k teammates as the key, and adopt the belief portrait representations z_k^t of the k teammates as the value. We then fuse the z_k^t , with the attention score represented as:

$$\alpha_{i,k} = \frac{\exp\left(\frac{1}{\sqrt{d_{key}}} (h_i^t W_Q) \cdot (\hat{h}_k^t W_K)^T\right)}{\sum_{j=1}^k \exp\left(\frac{1}{\sqrt{d_{key}}} (h_i^t W_Q) \cdot (\hat{h}_j^t W_K)^T\right)}, \quad (8)$$

where W_Q and W_K are learnable weight matrices applied to the query h_i^t and the key \hat{h}_k^t , while d_{key} denotes the dimension of the key vector. The attention score $\alpha_{i,k}$ represents the relevance between agent i and its collaborative teammates. The combined result is $e_i^t = \sum_{j=1}^k \alpha_{i,j} \cdot z_j^t$, which is concatenated with h_i^t and subsequently processed by a linear network to compute the local Q -value Q_i^{local} .

The optimization objective of this section is:

$$\mathcal{L}_{DF} = \lambda_{sy} \mathcal{L}_{sy} + \lambda_{se} \mathcal{L}_{se}, \quad (9)$$

where λ_{sy} and λ_{se} represent the hyperparameter of the loss function.

Overall Training Objective

Following the fusion of teammates' belief portraits, we utilize the basic CTDE training framework QMIX (Rashid et al. 2020) for training. Meanwhile, AIM remains compatible with other value decomposition methods, such as VDN (Sunehag et al. 2017) or QPLEX (Wang et al. 2020). All model parameters are updated by minimizing the \mathcal{L}_{TD} :

$$\mathcal{L}_{TD} = \mathbb{E} \left[(y - Q_{tot}(\tau, \mathbf{a}))^2 \right], \quad (10)$$

where $y = r + \gamma \max_{\mathbf{a}'} \hat{Q}_{tot}(\tau', \mathbf{a}')$ is the target network of the joint action-value function. By integrating the triple portrait and the dual filter, the complete training loss is as follows:

$$\mathcal{L}_{tot} = \mathcal{L}_{TD} + \mathcal{L}_{MD} + \mathcal{L}_{DF}. \quad (11)$$

Experiments

We select several methods as our primary baselines, including QMIX (Rashid et al. 2020), QPLEX (Wang et al. 2020), RODE (Wang et al. 2021), COLA (Xu et al. 2023), and SIRD (Zeng, Peng, and Li 2023). To evaluate the efficacy of our communication-free method in constructing teammates' decision logic, we include two communication-based methods: MAIC (Yuan et al. 2022) and T2MAC (Sun et al. 2024) as supplementary baselines. Furthermore, we set OMG (Yu, Jiang, and Lu 2024) with all teammates being trainable, serving as a baseline for agent modeling.




Figure 4: Ablation studies. (a)-(b) illustrate module-wise ablation. (c)-(d) present loss-wise ablation. “AIM_w/o_Belief” removes belief portrait, “AIM_w/o_Action” removes action portrait, while “AIM_w/o_Filter” removes dual filter.

Performance Comparisons

We conduct experiments on four MARL benchmarks: SMAC (Samvelyan et al. 2019), SMACv2 (Ellis et al. 2024), the Multi-agent Particle-Environment (MPE) (Mordatch and Abbeel 2018) and the Google Research Football (GRF) (Kurach et al. 2020). Notably, although GRF provides a fully observable environment, the soccer task demands frequent collaboration to achieve scoring goals, posing challenges to the agent modeling methods. For evaluation, we use five different random seeds and plot the average test results as bold lines. Due to page limitations, we leave the GRF experiments in Appendix D.1.

SMAC As the widely used benchmark in MARL, SMAC requires agents to master micro-level policies such as “focus fire,” “kite,” and “draw aggro” to defeat enemies under relatively concentrated initial positions. Figure 3 (a-f) presents the performance of AIM on six tasks, demonstrating optimal results on most maps. Particularly on maps that require a clear division of labor and the selection of the best collaborators, such as *3s5z_vs_3s6z*, *corridor*, and *6h_vs_8z*, AIM even outperforms communication-based baselines. This demonstrates that when agents are nearby, communication exerts negligible effects on decision-making quality. In contrast, AIM enables agents to select the most cooperative teammates by modeling teammates’ active inference processes. This capability enhances task allocation efficiency and contributes to achieving team objectives. In our OMG setup, each allied agent independently models teammates, failing to capture inter-agent coordination patterns and thus impairing collaboration.

SMACv2 As an extension of SMAC, SMACv2 introduces additional randomness in the initial positions and unit types, with dispersed agent distributions. It also employs true unit attack and sight ranges. The performance of AIM is shown in Figure 3 (g-l). Due to the difficulty in adapting to the environmental randomness, the baseline methods perform poorly. In contrast, AIM demonstrates superior adaptability to this uncertain environment. The results show that the proposed perception portrait module in AIM effectively adjusts to changes in the agents’ sight ranges, validating its efficacy. Furthermore, even with initially dispersed agent posi-

tions, AIM achieves performance comparable to or surpassing communication-based baselines. This validates the efficacy of AIM in assisting agents to understand teammates’ decisions and enhancing cooperation.

MPE The MPE provides a 2D physics-based environment supporting both continuous and discrete action spaces (Papoudakis et al. 2020). We evaluate the performance of AIM in the following three tasks with discrete actions: (1) Physical-Deception(PD)(2v1), (2) Predator-Prey(PP)(6v2), and (3) Predator-Prey(PP)(9v3). Since the original MPE is fully observable, we modify it to evaluate the effectiveness of the perception portrait. Specifically, we set the agents’ view radius to 0.8. As demonstrated in Figure 3 (m-o), AIM effectively infers teammates’ decision logic through modeling, enabling collaboration even under strict partial observability conditions.

Ablation Studies

In this section, we analyze the impact of active inference and dual filter modules on overall performance using ablation experiments. We progressively remove these key modules and evaluate their effects.

Active Inference Module We conduct ablation studies to validate the active inference module’s efficacy. Since the different sight ranges among units of SMACv2 have already validated the efficacy of the perception portrait, we focus on analyzing the belief and action portrait. We compare AIM with three ablation configurations: 1) AIM_w/o_Belief, which only removes the belief portrait; 2) AIM_w/o_Action, which only removes the action portrait; 3) AIM_w/o_Filter, which removes the dual filter of accuracy and relevance. The ablation results presented in Figure 4 (a)-(b) demonstrate that the removal of any individual module leads to significant performance degradation. This highlights that: (1) the comprehensive modeling of teammates’ decision processes is essential for maintaining accurate models, and (2) the dual filters are crucial for obtaining optimal teammates’ information.

To analyze the impact of different loss functions in AIM, we conduct ablation studies on each loss function. As demonstrated in Figures 4 (c)-(d), the elimination of any




Figure 5: Analysis of the parameter k in selective collaboration. Different values of k have varying impacts on performance.

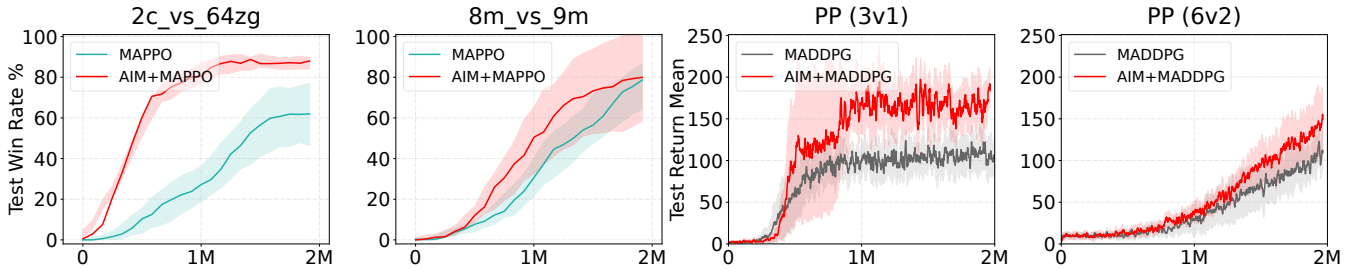


Figure 6: Scalability Experiments. AIM + MAPPO denotes the extension of AIM to the MAPPO framework, while AIM + MADDPG represents the adaptation of AIM to the MADDPG framework for continuous action spaces.

Maps	QMIX	AIM(Teammate)	AIM(Ours)
8m_vs_9m	87.2	89.1	94.0
3M_vs_2z5m	32.3	84.4	93.3
3s5z_vs_3s6z	00.0	50.4	58.7
6h_vs_8z	00.0	53.6	67.6

Table 1: Analysis of belief portraits from different perspectives. “AIM (Teammate)” models belief portraits from the teammates’ perspective, while AIM adopts the agent’s perspective for modeling. Bold indicates the best performance.

single loss function results in significant performance deterioration. Furthermore, the performance decline is more pronounced when \mathcal{L}_{cn} or \mathcal{L}_{se} is removed alone. This highlights the critical importance of maintaining continuous belief modeling of teammates while preserving self-focused decision cognition for enhancing complex task success rates.

Subsequently, we compare the construction of belief portraits from the agent’s perspective with that from the teammate’s perspective. Table 1 shows the AIM (Teammate) models belief portraits based on inaccurate perception portraits constructed from teammates’ perspectives. This approach introduces modeling biases that hinder collaboration. In contrast, AIM constructs belief portraits of teammates from the agent’s perspective, effectively reducing inaccuracies and improving cooperative performance.

Dual Filter Module We analyze the impact of the parameter k selected from the perception evaluation matrix, with a focus on its role in filtering redundant information. We conduct experiments using four super-hard maps from SMAC, and the results are shown in Figure 5. It is evident

that the choice of k significantly influences the performance of AIM. As k increases, collaboration with more teammates enhances task success rates due to increased information, but surpassing a certain threshold introduces redundancy, hindering decision-making. Therefore, the dual filter module optimizes the teammate selection, effectively preventing information redundancy and misleading decisions.

Scalability Across Policy-Based Frameworks

To further investigate the scalability of AIM across different policy-based frameworks, we extend AIM to both MAPPO (Yu et al. 2022) and MADDPG (Lowe et al. 2017). We conduct evaluations in both SMAC and continuous-action MPE. The results in Figure 6 demonstrate AIM + MAPPO surpassing MAPPO and AIM + MADDPG exceeding MADDPG. Experimental results validate the dual scalability capabilities of AIM: (1) extensibility within policy-based frameworks, and (2) adaptability to continuous action space environments.

Conclusion and Discussion

In this paper, we propose AIM, a novel framework that replaces explicit communication (i.e., “Tell”) with agent modeling (i.e., “Think”) to construct the cognition of teammates’ decision logic. AIM models teammates’ active inference through perception, belief, and action portraits, and employs a dual filter module to exclude inaccurate and irrelevant portraits. Experiments validate the effectiveness of AIM. However, in scenarios where selective collaboration is required, AIM selects a fixed set of top_k teammates as collaborators. The dynamic selection of collaborative teammates based on portraits remains an open issue and is left for future work.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62576029) and the Aeronautical Science Foundation of China (Grant No. 202300010M5001).

References

- Das, A.; Gervet, T.; Romoff, J.; Batra, D.; Parikh, D.; Rabbat, M.; and Pineau, J. 2019. Tarmac: Targeted multi-agent communication. In *International Conference on machine learning*, 1538–1546. PMLR.
- Ding, Z.; Huang, T.; and Lu, Z. 2020. Learning individually inferred communication for multi-agent cooperation. *Advances in Neural Information Processing Systems*, 33: 22069–22079.
- Ellis, B.; Cook, J.; Moalla, S.; Samvelyan, M.; Sun, M.; Mahajan, A.; Foerster, J.; and Whiteson, S. 2024. Smacv2: An improved benchmark for cooperative multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; Pezzulo, G.; et al. 2016. Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68: 862–879.
- He, H.; Boyd-Graber, J.; Kwok, K.; and Daumé III, H. 2016. Opponent modeling in deep reinforcement learning. In *International conference on machine learning*, 1804–1813. PMLR.
- Hu, G.; Zhu, Y.; Zhao, D.; Zhao, M.; and Hao, J. 2023. Event-Triggered Communication Network With Limited-Bandwidth Constraint for Multi-Agent Reinforcement Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8): 3966–3978.
- Kiran, B. R.; Sobh, I.; Talpaert, V.; Mannion, P.; Al Salhab, A. A.; Yogamani, S.; and Pérez, P. 2021. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6): 4909–4926.
- Kurach, K.; Raichuk, A.; Stańczyk, P.; Zajac, M.; Bachem, O.; Espenholt, L.; Riquelme, C.; Vincent, D.; Michalski, M.; Bousquet, O.; et al. 2020. Google research football: A novel reinforcement learning environment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 4501–4510.
- Lee, J.; Chung, J.; and Sohn, K. 2019. Reinforcement learning for joint control of traffic signals in a transportation network. *IEEE Transactions on Vehicular Technology*, 69(2): 1375–1387.
- Li, D.; Lou, N.; Xu, Z.; Zhang, B.; and Fan, G. 2025. Efficient Communication in Multi-Agent Reinforcement Learning with Implicit Consensus Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 23240–23248.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Mordatch, I.; and Abbeel, P. 2018. Emergence of grounded compositional language in multi-agent populations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Oliehoek, F. A.; Amato, C.; et al. 2016. *A concise introduction to decentralized POMDPs*, volume 1. Springer.
- Papoudakis, G.; and Albrecht, S. V. 2020. Variational Autoencoders for Opponent Modeling in Multi-Agent Systems. *CoRR*, abs/2001.10829.
- Papoudakis, G.; Christianos, F.; and Albrecht, S. 2021. Agent modelling under partial observability for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 19210–19222.
- Papoudakis, G.; Christianos, F.; Schäfer, L.; and Albrecht, S. V. 2020. Benchmarking multi-agent deep reinforcement learning algorithms in cooperative tasks. *arXiv preprint arXiv:2006.07869*.
- Rabinowitz, N.; Perbet, F.; Song, F.; Zhang, C.; Eslami, S. A.; and Botvinick, M. 2018. Machine theory of mind. In *International conference on machine learning*, 4218–4227. PMLR.
- Raileanu, R.; Denton, E.; Szlam, A.; and Fergus, R. 2018. Modeling others using oneself in multi-agent reinforcement learning. In *International conference on machine learning*, 4257–4266. PMLR.
- Rashid, T.; Samvelyan, M.; De Witt, C. S.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research*, 21(178): 1–51.
- Roesch, M.; Linder, C.; Zimmermann, R.; Rudolf, A.; Hohmann, A.; and Reinhart, G. 2020. Smart grid for industry using multi-agent reinforcement learning. *Applied Sciences*, 10(19): 6900.
- Samvelyan, M.; Rashid, T.; De Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*.
- Song, S.; Lin, Y.; Han, S.; Yao, C.; Wu, H.; Wang, S.; and Lv, K. 2025. CoDe: Communication Delay-Tolerant Multi-Agent Collaboration via Dual Alignment of Intent and Timeliness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(22): 23304–23312.
- Sun, C.; Zang, Z.; Li, J.; Li, J.; Xu, X.; Wang, R.; and Zheng, C. 2024. T2MAC: Targeted and Trusted Multi-Agent Communication through Selective Engagement and Evidence-Driven Integration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15154–15163.
- Sun, X. 2024. *Assessing Model Robustness in Complex Visual Environments*. Ph.D. thesis, The Australian National University (Australia).
- Sun, X.; Leng, X.; Wang, Z.; Yang, Y.; Huang, Z.; and Zheng, L. 2023a. Cifar-10-warehouse: Broad and more realistic testbeds in model generalization analysis. *arXiv preprint arXiv:2310.04414*.
- Sun, X.; Yao, Y.; Wang, S.; Li, H.; and Zheng, L. 2023b. Alice Benchmarks: Connecting Real World Re-Identification with the Synthetic. *arXiv preprint arXiv:2310.04416*.

- Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W. M.; Zambaldi, V.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; et al. 2017. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296*.
- Tian, Z.; Wen, Y.; Gong, Z.; Punakkath, F.; Zou, S.; and Wang, J. 2019. A regularized opponent model with maximum entropy objective. *arXiv preprint arXiv:1905.08087*.
- Tung, T.-Y.; Kobus, S.; Roig, J. P.; and Gündüz, D. 2021. Effective Communications: A Joint Learning and Communication Framework for Multi-Agent Reinforcement Learning Over Noisy Channels. *IEEE Journal on Selected Areas in Communications*, 39(8): 2590–2603.
- Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2020. Qplex: Duplex dueling multi-agent q-learning. *arXiv preprint arXiv:2008.01062*.
- Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2021. {RODE}: Learning Roles to Decompose Multi-Agent Tasks. In *International Conference on Learning Representations*.
- Wang, T.; Wang, J.; Zheng, C.; and Zhang, C. 2019. Learning nearly decomposable value functions via communication minimization. *arXiv preprint arXiv:1910.05366*.
- Xie, A.; Losey, D.; Tolsma, R.; Finn, C.; and Sadigh, D. 2021. Learning latent representations to influence multi-agent interaction. In *Conference on robot learning*, 575–588. PMLR.
- Xu, Z.; Zhang, B.; Li, D.; Zhang, Z.; Zhou, G.; Chen, H.; and Fan, G. 2023. Consensus learning for cooperative multi-agent reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11726–11734.
- Xue, W.; Qiu, W.; An, B.; Rabinovich, Z.; Obraztsova, S.; and Yeo, C. K. 2021. Mis-spoke or mis-lead: Achieving robustness in multi-agent communicative reinforcement learning. *arXiv preprint arXiv:2108.03803*.
- Yang, T.; Meng, Z.; Hao, J.; Zhang, C.; Zheng, Y.; and Zheng, Z. 2018. Towards efficient detection and optimal response against sophisticated opponents. *arXiv preprint arXiv:1809.04240*.
- Yao, C.; Lin, Y.; Song, S.; Wu, H.; Ma, Y.; Han, S.; and Lv, K. 2025. From General Relation Patterns to Task-Specific Decision-Making in Continual Multi-Agent Coordination. *arXiv preprint arXiv:2507.06004*.
- Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35: 24611–24624.
- Yu, X.; Jiang, J.; and Lu, Z. 2024. Opponent modeling based on subgoal inference. *Advances in Neural Information Processing Systems*, 37: 60531–60555.
- Yuan, L.; Wang, J.; Zhang, F.; Wang, C.; Zhang, Z.; Yu, Y.; and Zhang, C. 2022. Multi-agent incentive communication via decentralized teammate modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 9466–9474.
- Zeng, X.; Peng, H.; and Li, A. 2023. Effective and stable role-based multi-agent collaboration by structural information principles. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 11772–11780.
- Zhai, Y.; Peng, P.; Su, C.; and Tian, Y. 2023. Dynamic Belief for Decentralized Multi-Agent Cooperative Learning. In Elkind, E., ed., *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 344–352.
- Zhu, C.; Dastani, M.; and Wang, S. 2024. A survey of multi-agent deep reinforcement learning with communication. *Autonomous Agents and Multi-Agent Systems*, 38(1): 4.
- Zintgraf, L.; Devlin, S.; Ciosek, K.; Whiteson, S.; and Hofmann, K. 2021. Deep interactive bayesian reinforcement learning via meta-learning. *arXiv preprint arXiv:2101.03864*.