

Contrastive Identity-Aware Learning for Multi-Agent Value Decomposition

Shunyu Liu*, Yihe Zhou*, Jie Song†, Tongya Zheng, Kaixuan Chen,
Tongtian Zhu, Zunlei Feng, Mingli Song

Zhejiang University

liushunyu@zju.edu.cn, yihe_zhou@zju.edu.cn, sjie@zju.edu.cn, tyzheng@zju.edu.cn, chenkx@zju.edu.cn,
raiden@zju.edu.cn, zunleifeng@zju.edu.cn, brooksong@zju.edu.cn

Abstract

Value Decomposition (VD) aims to deduce the contributions of agents for decentralized policies in the presence of only global rewards, and has recently emerged as a powerful credit assignment paradigm for tackling cooperative Multi-Agent Reinforcement Learning (MARL) problems. One of the main challenges in VD is to promote diverse behaviors among agents, while existing methods directly encourage the diversity of learned agent networks with various strategies. However, we argue that these dedicated designs for agent networks are still limited by the indistinguishable VD network, leading to homogeneous agent behaviors and thus downgrading the cooperation capability. In this paper, we propose a novel Contrastive Identity-Aware learning (CIA) method, explicitly boosting the credit-level distinguishability of the VD network to break the bottleneck of multi-agent diversity. Specifically, our approach leverages contrastive learning to maximize the mutual information between the temporal credits and identity representations of different agents, encouraging the full expressiveness of credit assignment and further the emergence of individualities. The algorithm implementation of the proposed CIA module is simple yet effective that can be readily incorporated into various VD architectures. Experiments on the SMAC benchmarks and across different VD backbones demonstrate that the proposed method yields results superior to the state-of-the-art counterparts. Our code is available at <https://github.com/liushunyu/CIA>.

Introduction

Cooperative Multi-agent Reinforcement Learning (MARL) aims to jointly train multiple agents to achieve one common goal, and has witnessed its unprecedented success in various applications, such as video games (Vinyals et al. 2019; Jaderberg et al. 2019), traffic light systems (Wu et al. 2020), and smart grid control (Yan and Xu 2020; Xu et al. 2020; Wang et al. 2021b). To alleviate the partial observation constraint and the scalability problem in cooperative MARL, the *Centralized Training and Decentralized Execution* (CTDE) framework has attracted increasing attention (Lowe et al. 2017), where agents are granted access to additional global information during centralized training and deliver actions

*These authors contributed equally.

†Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

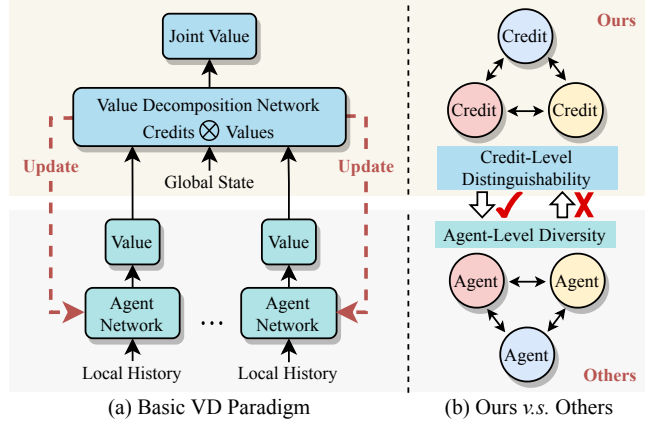


Figure 1: (a) Basic VD paradigm. \otimes denotes the complex network operator based on different VD methods, where credits can be considered as the agent contributions. (b) Comparing the credit-level distinguishability mechanism with the agent-level diversity mechanism under the VD paradigm. Credit-level distinguishability can promote agent-level diversity, but not vice versa.

only based on their local histories in a decentralized way. The decentralized policies with one shared network significantly reduce the number of learnable parameters and the exponentially growing joint-action space to the linear complexity. Nevertheless, the CTDE framework still faces a critical challenge of credit assignment, *i.e.*, how to deduce the contributions of individual agents given only global rewards.

In recent years, Value Decomposition (VD) has emerged as a promising credit assignment paradigm, which endows agents with the ability to learn their own optimal policies by decomposing the joint value function according to their individual credits, as depicted in Figure 1(a). The seminal work, VDN (Sunehag et al. 2018), additively factorizes the team value function into individual agent-wise terms, imposing a strong constraint of credit equalization. QMIX (Rashid et al. 2018) identifies that the full factorization of VDN is not necessary and relaxes this additive constraint to a non-linear combination, which enables the weighted credits for per-agent values. Due to the superior performance of QMIX, there have been many recent efforts to improve its represen-

tation capability (Rashid et al. 2020; Wang et al. 2021a).

However, despite the large representation capacity of existing VD methods, it is widely observed that agents often learn similar behaviors inevitably (Mahajan et al. 2019; Hu et al. 2022). Such similar behaviors easily lead to the local optimum of the cooperative policies, which may severely impede the efficient exploration and downgrade the final performance (Terry et al. 2020). Consider a football game where different agents observe similarly. If agents behave similarly for desired credits, all of them may gather to compete for a ball. However, they should keep distance and adopt diverse tactics. Several studies attribute this similarity between agents to the homogeneous policy networks with parameter sharing (Jiang and Lu 2021; Li et al. 2021). They propose to design auxiliary objectives in policy networks towards agent-level diversity, and even introduce agent-specific networks for each agent while sacrificing the advantage of complete parameter sharing. However, these works ignore the fact that policy networks are evaluated and improved via the VD network, where the diversity of agent behaviors often depends on the distinguishability of credit assignment, as depicted in Figure 1(b). Thus, we argue that their pursuit of diverse policies is still limited when the VD network provides ambiguous credits to demonstrate the contributions of different agents.

In this paper, we investigate the multi-agent diversity from a novel perspective of credit assignment. To evaluate the distinguishability of existing VD methods, we design a random-shuffle scheme that eliminates the identity information of input agent values while preserving the original network architecture. Case studies empirically show that the VD network may assign ambiguous credits to the agents, indicating that promoting the diversity of agents will be limited by the identity insensitivity of the VD network. Therefore, we introduce a new contrastive identity-aware learning method, termed as CIA, to explicitly encourage credit-level distinguishability. The proposed method leverages gradient-based attribution to represent the credit of each agent, and further considers the long-term behaviors of agents by adopting the temporal credits from the overall trajectory. To encourage the discriminative assignment of credits and further the emergence of diverse behaviors, we propose to maximize the mutual information between the temporal credits and learnable identity representations of different agents. Moreover, we customize a contrastive learning objective to derive a tractable lower bound for the mutual information since estimating and maximizing the mutual information of neural networks is often intractable. Our main contributions can be summarized as follows:

- We identify the ambiguous credit assignment problem in VD, a highly important ingredient for multi-agent diversity yet largely overlooked by existing literature.
- We propose a novel contrastive identity-aware learning (CIA) method to promote diverse behaviors via explicitly encouraging credit-level distinguishability. The proposed CIA module imposes no constraints over the network architecture, and serves as a plug-and-play module readily applicable to various VD methods.

- Experiments conducted on the StarCraft II micromanagement benchmark show that CIA yields significantly superior performance to state-of-the-art competitors.

Related Work

Value Decomposition aims to extract the individual utility from the global reward for credit assignment, which has become a well-established paradigm for tackling cooperative MARL problems (Zohar, Mannor, and Tennenholtz 2022; Jeon et al. 2022; Fu et al. 2022). To realize efficient VD, it is critical to satisfy the Individual-Global-Max (IGM) principle that the global optimal action should be consistent with the collection of individual optimal actions of agents (Son et al. 2019). Following the IGM principle, VDN (Sunehag et al. 2018) proposes to represent the joint value function as a sum of individual value functions, while QMIX (Rashid et al. 2018) extends this additive VD and imposes a monotonicity constraint, showing the state-of-the-art performance. To further alleviate the risk of suboptimal results, WQMIX (Rashid et al. 2020) improves QMIX by a weighted projection that allows more emphasis to be placed on underestimated actions.

However, VDN and QMIX still suffer from structural constraints which limit their representation capability for joint action-value function classes. Thus, QTRAN (Son et al. 2019) constructs the soft regularization constraints to achieve more general decomposition than VDN and QMIX. QPLEX (Wang et al. 2021a), on the other hand, proposes a duplex dueling network architecture to enable the complete VD function class that satisfies the IGM principle. These VD methods have shown great potential in the field of credit assignment for challenging MARL tasks. Other works further adapt VD towards transfer learning (Long et al. 2020; Wang et al. 2020b; Yang et al. 2022; Yang, Ye, and Wang 2022) and ad hoc teamwork (Gu et al. 2021; Macke, Mirsky, and Stone 2021; Rahman et al. 2021). Despite the success of VD, existing methods ignore the distinguishability in credit assignment, leading to the homogeneous behaviors among multiple agents. GRE (Zhao et al. 2022) tries to use credit entropy regularization to mitigate this problem, which only imposes single-step distinguishability on credits.

Agent Diversity has been widely studied in single-agent reinforcement learning problems, which provides an exploration bonus to encourage the diverse behaviors (Eysenbach et al. 2019; Burda et al. 2019). In recent years, due to the promising results achieved by the existing single-agent methods, diversity has also emerged as a popular topic in MARL (Lee, Yang, and Lim 2019; Christianos et al. 2021). It encourages the difference between individual agents under the cooperative setting when pursuing diversity in the context of MARL. RODE (Wang et al. 2021c) realizes a learnable role assignment for agents to achieve diversity. It learns action representations and clusters different actions into several restricted roles, while this role-based method is limited by the action space that cannot be decomposed. EOI (Jiang and Lu 2021) proposes to construct the intrinsic reward, a predicted probability of agent identity given its observation, to promote the emergence of individuality.

However, the identity prediction mechanism of EOI is easy to overfit if the local observation contains the identity information. In MAVEN (Mahajan et al. 2019), agents condition their behaviors on the shared latent variable controlled by a hierarchical policy, where the mutual information between the trajectories and latent variables is maximized to learn a diverse set of such behaviors. Similarly, EITI (Wang et al. 2020a) leverages mutual information to capture the influence between the transition dynamics of different agents, while CDS (Li et al. 2021) designs agent-specific networks for each agent and encourages diversity by optimizing the mutual information between the agent identities and trajectories. These works are all built based on the existing VD methods, especially QMIX, encouraging only the diversity of agent networks with various strategies. Thus, they still suffer from limited distinguishability when credit assignment in the VD network is ambiguous.

Preliminary

Dec-POMDP. We consider a fully cooperative multi-agent setting under the *Decentralized Partially Observable Markov Decision Process* (Dec-POMDP), which is defined as a tuple $\langle \mathcal{A}, \mathcal{S}, \mathcal{U}, P, r, \Omega, O, \gamma \rangle$, where $\mathcal{A} = \{a_k\}_{k=1}^K$ is the set of K agents and $s \in \mathcal{S}$ is the global state of the environment. At each time step t , each agent $a_k \in \mathcal{A}$ receives an individual partial observation $o_t^k \in \Omega$ drawn from the observation function $O(s_t, a_k)$. Then each agent chooses a corresponding action $u_t^k \in \mathcal{U}$, forming a joint action $\mathbf{u}_t \in \mathcal{U}^K$. This causes a transition to the next state s_{t+1} according to the state transition function $P(s_{t+1}|s_t, \mathbf{u}_t) : \mathcal{S} \times \mathcal{U}^K \times \mathcal{S} \rightarrow [0, 1]$. All agents share the same reward function $r(s_t, \mathbf{u}_t) : \mathcal{S} \times \mathcal{U}^K \rightarrow \mathbb{R}$ and $\gamma \in [0, 1]$ is the discount factor. Each agent a_k has an action-observation history $\tau^k \in \mathcal{T} \equiv (\Omega \times \mathcal{U})^*$ and learns its individual policy $\pi^k(u^k|\tau^k) : \mathcal{T} \times \mathcal{U} \rightarrow [0, 1]$ to jointly maximize the discounted return $R_t = \sum_{i=0}^{\infty} \gamma^i r_{t+i}$. The joint action-observation history is define as $\tau \in \mathcal{T}^K$. The joint policy π induces a joint action-value function $Q_t^{tot}(s_t, \mathbf{u}_t) = \mathbb{E}_{s_{t+1:\infty}, \mathbf{u}_{t+1:\infty}} [R_t | s_t, \mathbf{u}_t, \pi]$ that represents the expected discounted return under the given policy.

The CTDE Framework has attracted substantial attention in cooperative MARL to achieve effective policy learning (Yu et al. 2021; Luo et al. 2022), where agents must learn the decentralized policies which based on only local observation at execution time, but they are granted access to the global information during centralized training. One of the promising ways to exploit the CTDE framework is VD, which allows agents to learn their individual utility functions by optimizing the joint action-value function for credit assignment. To realize VD, a mixing network with parameters θ_v is adopted as an approximator to estimate the joint action-value function Q^{tot} . The mixing network is introduced to merge all individual action values into a joint one $Q^{tot} = f(\mathbf{q}; \theta_v)$, where $\mathbf{q} = [Q^k]_{k=1}^K \in \mathbb{R}^K$ and Q^k with shared parameters θ_π is the utility network of each agent a_k . The learnable parameter $\theta = \{\theta_\pi, \theta_v\}$ can be updated by

minimizing the following Temporal-Difference (TD) loss:

$$\mathcal{L}_{TD}(\theta) = \mathbb{E}_{\mathcal{D}} \left[(y^{tot} - Q^{tot})^2 \right]. \quad (1)$$

where $\mathbb{E}[\cdot]$ denotes the expectation function, \mathcal{D} is the replay buffer of the transitions, $y^{tot} = r + \gamma \hat{Q}^{tot}$ is the one-step target and \hat{Q}^{tot} is the target network (Mnih et al. 2015).

Method

In what follows, we provide case studies to evaluate the credit indistinguishability of existing VD methods. Then we further detail the proposed contrastive identity-aware learning (CIA) module and finally summarize the complete optimization algorithm.

Credit Indistinguishability Analysis

Firstly, we define that a credit assignment is ambiguous if the learnable credits are invariant to the agent identities. To investigate the ambiguous credit assignment problem in VD, we design a training scheme that randomly shuffles the order of input values to the mixing network at every training epoch. Concretely, we denote Φ as the set of all permutation matrices $\mathbf{P} \in \mathbb{R}^{K \times K}$. At each training epoch, we randomly sample a permutation matrix $\mathbf{P} \in \Phi$ and obtain the ambiguous joint-action value $Q^{tot} = f(\mathbf{P}\mathbf{q})$ instead of the original one to calculate the TD loss and update the network parameter. This random-shuffle scheme eliminates the identity information of input agent values without changing the original network architecture, providing a comparable baseline with ambiguous credit assignment. Furthermore, we propose to use KL-divergence distance to measure the similarity of

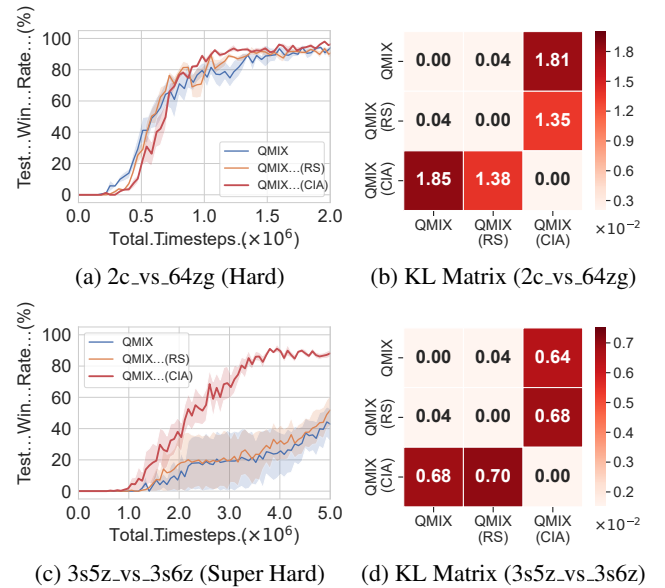


Figure 2: (Left) Learning curves of QMIX and its variants on two SMAC scenarios. (Right) KL-divergence distance matrices of the credit distributions of different methods.

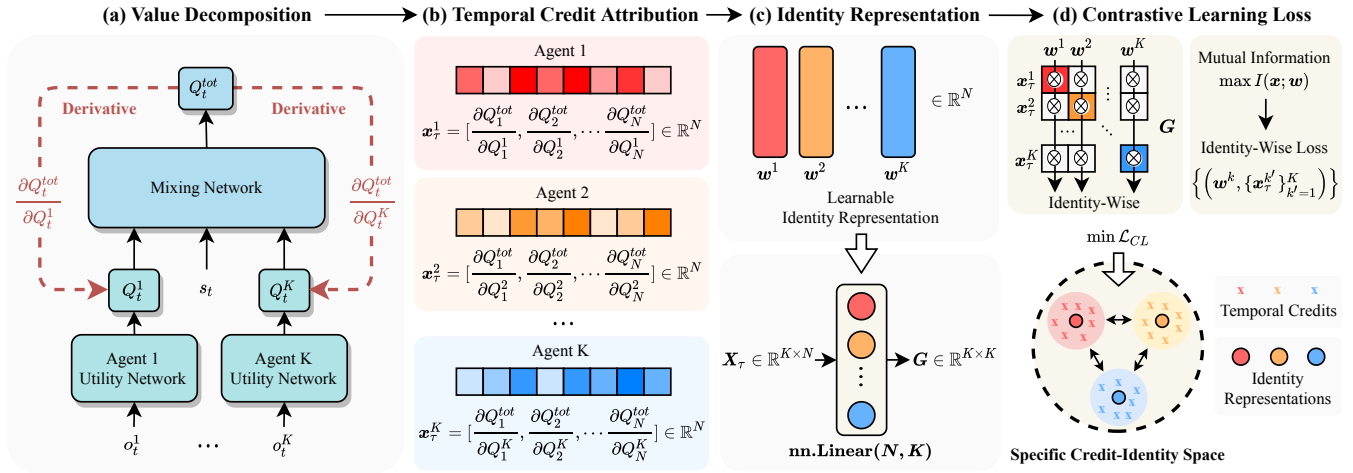


Figure 3: An illustrative diagram of the proposed contrastive identity-aware learning (CIA) method.

credit distributions¹ of different methods. We calculate the average KL-divergence distance by the trajectories sampled by the trained networks of all compared methods separately. Please refer to Appendix B for more details about the experiment settings.

Case studies are conducted on two challenging tasks provided by SMAC (Samvelyan et al. 2019). We take QMIX (Rashid et al. 2018) as an example to study due to its state-of-the-art performance, and introduce its variant under the random-shuffle training scheme, namely QMIX (RS). Moreover, we also integrate the proposed CIA module with QMIX, namely QMIX (CIA), to demonstrate the importance of credit-level distinguishability. Experimental results of different methods are shown in Figure 2. Interestingly, even with the ambiguous credit assignment, the random-shuffle variant can still achieve the performance on par with those obtained by original QMIX. The KL-divergence distance also shows that the credit distribution of QMIX is similar as its random-shuffle variant, indicating that QMIX are insensitive to the identity of agents. Thus, its learned credits may be ambiguous, which limits the diverse behaviors of agents and damages the final performance.

Contrastive Identity-Aware Learning

To encourage multi-agent diversity, we introduce a novel contrastive learning method to realize the identity-aware distinguishability in credit assignment, as shown in Figure 3.

Temporal Credit Attribution. Before promoting the distinguishability of credits, one key issue to be considered is: how to define a universal credit to represent the contributions of individual agents in different VD methods. Towards addressing this problem, we consider a gradient-based attribution mechanism to extract the credit assignment information in the mixing network. In general, the gradient-based attribution is calculated taking the partial derivatives of the output

¹At each time step t , the credit distribution \mathbf{d}_t is defined as the normalized gradient-based credits (mentioned in Eq. 2) over all agents: $\mathbf{d}_t = \text{softmax}([x_t^1, x_t^2, \dots, x_t^K])$.

with respect to the input, indicating the relative importance of different inputs to a specific output (Ancona et al. 2018). Thus, we employ this mechanism to attribute the credits as the partial derivatives between the joint-action value and individual values, as shown in Figure 3(a). Specially, the credit of the agent a_k at each time step t is formulated as:

$$x_t^k = \frac{\partial Q_t^{tot}}{\partial Q_t^k} \in \mathbb{R}. \quad (2)$$

This credit attribution mechanism enables us to determine the contribution of each agent value, also known as the value sensitivity. That is to say, a larger attribution magnitude means that the corresponding agent has a significant impact on the final result. Furthermore, this credit attribution mechanism can be readily applied to different VD methods, regardless of the heterogeneous mixing network architectures.

In sequential decision-making problem, it is insufficient to deduce the contribution of an agent only by the behavior of a single step. The behavior of an agent is usually affected by its final goal. As shown in Figure 3(b), we therefore adopt the temporal credit attribution of each agent a_k from the overall sampled trajectory τ with $t \in \{1, 2, \dots, N\}$ as:

$$\mathbf{x}_\tau^k = [x_1^k, x_2^k, \dots, x_N^k] \in \mathbb{R}^N, \quad (3)$$

where N is the maximum length of the sampled trajectories. If any trajectories are truncated early due to the game rules, zero padding will be applied to align the lengths of all trajectories. We denote $\mathbf{X}_\tau \in \mathbb{R}^{K \times N}$ as the temporal credits of K agents in the trajectory τ . This temporal credit attribution considers the long-term behaviors of agents, leading to a more stable measure for the agent contributions.

Contrastive Identity-Aware Distinguishability. After defining the credits in a reasonable way, we further attempt to promote the distinguishability of the credits, avoiding capacity degradation of the VD model due to the ambiguous credit assignment problem. However, obtaining the solution that all the credits differ from each other is not trivial. It is hard to directly constrain the distance between different

learnable credits during the optimization process. Thus, we approximate the solution by introducing latent identity representations for each agent as intermediate variables, and condition the individual temporal credits on these representations, realizing identity-aware distinguishability.

Specially, each agent a_k is assigned only one learnable random-initialized identity representation $\mathbf{w}^k \in \mathbb{R}^N$ with the same dimension of N as the temporal credits during the entire training process. We denote $\mathbf{W} \in \mathbb{R}^{K \times N}$ as the identity representations of K agents. On the other hand, it is notable that the temporal credits \mathbf{x}_τ^k of the agent are various in each sampled trajectory τ . We propose to maximize the mutual information $I(\mathbf{x}; \mathbf{w})$ between the temporal credits and identity representations of different agents:

$$I(\mathbf{x}; \mathbf{w}) = \mathbb{E}_{\mathbf{x}, \mathbf{w}} \left[\log \frac{p(\mathbf{x} | \mathbf{w})}{p(\mathbf{x})} \right]. \quad (4)$$

However, directly optimizing this objective is quite difficult since estimating mutual information is often intractable. Inspired by contrastive learning (Oord, Li, and Vinyals 2018), we introduce a contrastive learning loss, the InfoNCE loss, to provide a traceable lower bound of the mutual information objective as follows:

$$I(\mathbf{x}; \mathbf{w}) \geq \log(K) - \mathcal{L}_{CL}, \quad (5)$$

where \mathcal{L}_{CL} is the InfoNCE loss and K is the number of agents. Contrastive learning can be considered as learning a differentiable dictionary look-up task, which contrasts semantically similar and dissimilar query-key pairs. To match individual temporal credits with their corresponding identity representations, we define identity representations as queries and temporal credits as keys. For example, given a query \mathbf{w}^k and keys $\mathcal{X} = \{\mathbf{x}_\tau^{k'}\}_{k'=1}^K$, the goal of contrastive learning is to ensure that \mathbf{w}^k is close with \mathbf{x}_τ^k while being irrelevant to other keys in $\mathcal{X} \setminus \{\mathbf{x}_\tau^k\}$. Specially, the identity-wise contrastive learning loss is given as follows:

$$\mathcal{L}_{CL} = \mathbb{E}_{(\mathbf{w}^k, \{\mathbf{x}_\tau^{k'}\}_{k'=1}^K) \sim \mathcal{D}} \left[-\log \frac{\exp(g(\mathbf{x}_\tau^k, \mathbf{w}^k))}{\sum_{k'=1}^K \exp(g(\mathbf{x}_\tau^{k'}, \mathbf{w}^k))} \right], \quad (6)$$

where $g(\mathbf{x}_\tau, \mathbf{w}) = \mathbf{x}_\tau^T \mathbf{w} \in \mathbb{R}$ is the function that compute the non-negative similarity score between the identity representation \mathbf{w} and the temporal credit \mathbf{x}_τ . Here we adopt a simple dot-product similarity, but it can be easily replaced by other methods. The similarity matrix is denoted as $\mathbf{G} = \mathbf{X}_\tau \mathbf{W}^T \in \mathbb{R}^{K \times K}$. As shown in Figure 3(d), this identity-wise contrastive learning loss constrains the learned identity representations to uniformly distribute on a specific credit-identity hypersphere without divergence, where the temporal credits distribute around their corresponding identity representation (Wang and Isola 2020). With the contrastive learning loss, CIA directly encourages the identity-aware distinguishability among agent credits, providing a more discriminative credit assignment for multi-agent diversity.

Intuitively, we can explain the intrinsic mechanism in CIA from the perspective of a classification problem, where each input credits predict their corresponding identity labels. The

identity representations are the weights of the classifier. For classification, there are two ways to obtain a minimal loss: use some easier inputs or adopt a more complex classifier. Here the identity representations only form a simple one-layer linear classifier, which imposes a strong constraint that the input credits must be linearly separable for minimal loss. Thus, the CIA module successfully endows the VD network the ability to promote credit distinguishability.

Optimization Algorithm

To sum up, training our framework based on the CIA module contains two main parts. The first one is the original TD loss \mathcal{L}_{TD} , which enables each agent to learn its individual policy by optimizing the joint-action value of the mixing network. The second one is the proposed contrastive learning loss \mathcal{L}_{CL} to facilitate the credit-level distinguishability. Thus, given these two corresponding loss items, the total loss of our framework is formulated as follows:

$$\mathcal{L}_{all} = \mathcal{L}_{TD} + \alpha \mathcal{L}_{CL}, \quad (7)$$

where α is the coefficient for trading off loss terms. The overall framework is trained in an end-to-end centralized manner. It is notable that CIA does not break any constraints of original VD paradigm and still follows IGM principle. To make the proposed CIA clearer to readers, we provide the pseudocode in Appendix A. We only need to introduce an additional linear layer for learnable identity representations as shown in Figure 3(c). The algorithm implementation of the CIA module is simple yet effective that can be seamlessly integrated with various VD architectures.

Experiments

To demonstrate the effectiveness of the proposed CIA method, we conduct experiments on the didactic game and the StarCraft II micromanagement challenge.

Experimental Settings

Our methods are compared with various state-of-the-art methods, including (i) Value decomposition methods: QMIX (Rashid et al. 2018), QPLEX (Wang et al. 2021a), QTRAN (Son et al. 2019), OWQMIX and CWQMIX (Rashid et al. 2020). (ii) Diversity-based methods: MAVEN (Mahajan et al. 2019), EOI (Jiang and Lu 2021), GRE (Zhao et al. 2022), SCDS that is a variant of CDS (Li et al. 2021) with complete shared agent network to ensure comparability. Our CIA implementation uses QMIX and QPLEX as the integrated backbones to evaluate its performance, namely QMIX (CIA) and QPLEX (CIA). These two methods are chosen for their robust performance in different multi-agent tasks, while CIA can also be readily applied to other methods. We adopt the Python MARL framework (PyMARL) (Samvelyan et al. 2019) to implement our method and all baselines. The detailed hyperparameters are given in Appendix B, where the common training parameters across different methods are consistent.

Didactic Game

We design a didactic game and call it *Turn*. Two colored agents with the observation of 3×3 are initialized in a 5×5

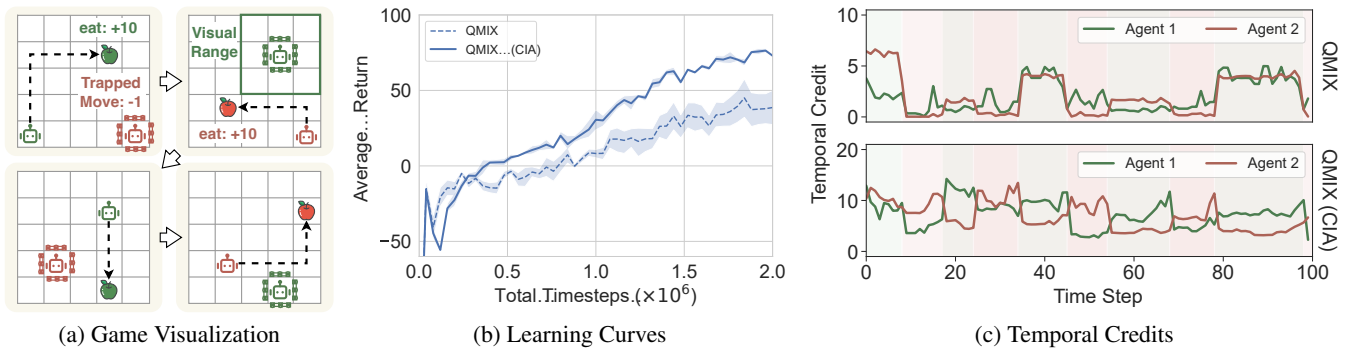


Figure 4: The performance comparison on the didactic game of Turn.

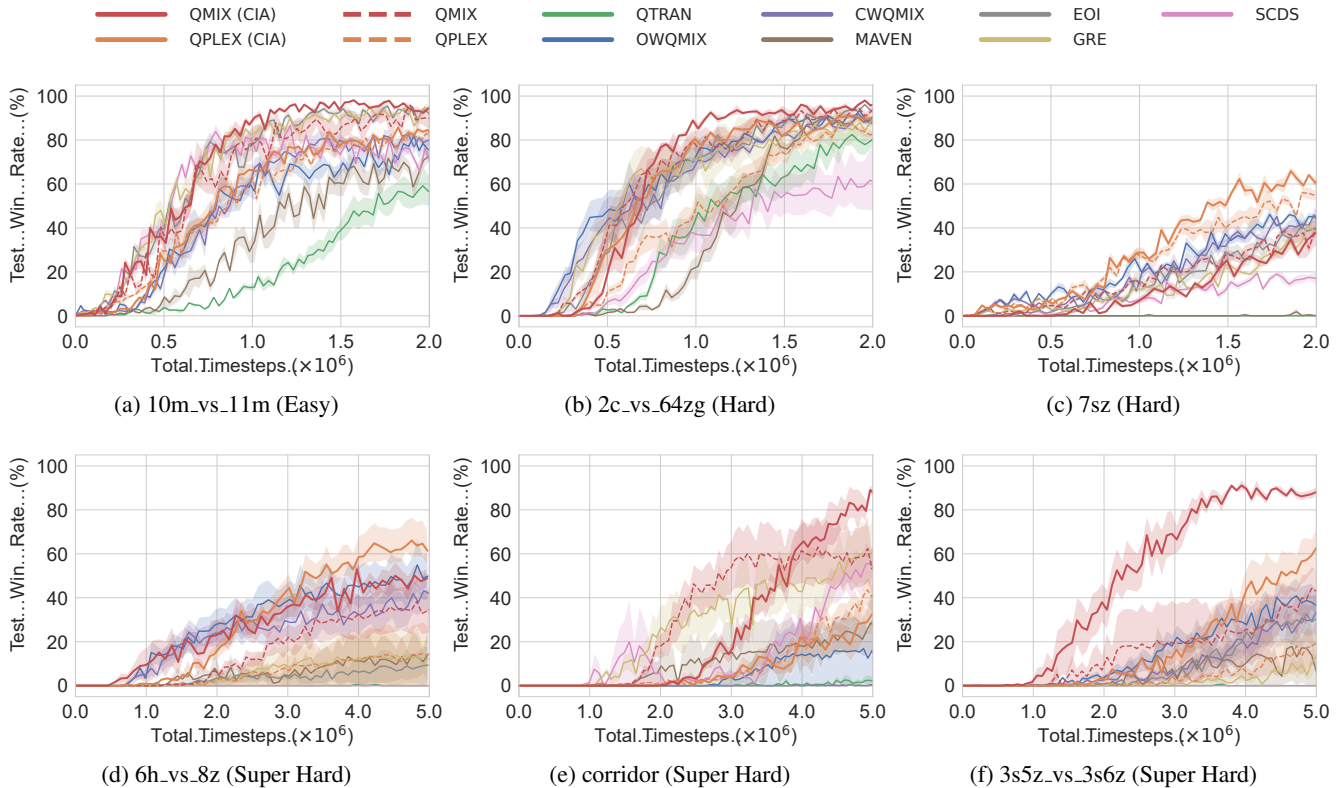


Figure 5: Learning curves of our proposed CIA variants and baselines on the SMAC scenarios. All experimental results are illustrated with the mean and the standard deviation of the performance over five random seeds for a fair comparison. To make the results in figures clearer for readers, we adopt a 50% confidence interval to plot the error region.

map, and the corresponding colored apples are generated alternately in the map. The goal of each agent is to take turns to eat the exclusive apples (agent 1 first), where the one that is not in its round will be trapped by fences and punished when moving. This didactic game requires the VD model to learn a rotating credit assignment strategy. As shown in Figure 4, the proposed CIA works quite well on improving QMIX. We further visualize the temporal credits of different methods in a sampled trajectory. Obviously, CIA realizes the distinguishable credit assessment, where the temporal credits of two agents rise alternately. However, the credits of QMIX

are sometimes ambiguous, which even pay more attention to the trapped agent in each round. This crisp example verifies the effectiveness of the proposed CIA. Please refer to Appendix B for more details.

SMAC Benchmark

The StarCraft Multi-Agent Challenge (SMAC)² (Samvelyan et al. 2019) has become a common-used benchmark for evaluating state-of-the-art MARL methods. SMAC focuses on

²We use SC2.4.10 version instead of the older SC2.4.6.2.69232. Performance is not comparable across versions.

micromanagement challenges where each of the ally entities is controlled by an individual learning agent, and enemy entities are controlled by a built-in AI. The goal of the agents is to maximize the damage to enemies. Hence, proper tactics such as sneaking attack and drawing fire are required during battles. Learning these diverse interaction behaviors under partial observation is a challenging task. To validate the effectiveness of our methods, we conduct experiments on 6 SMAC scenarios (Samvelyan et al. 2019) which are classified into **Easy** (*10m_vs_11m*), **Hard** (*2c_vs_64zg*, *7sz*) and **Super Hard** (*6h_vs_8z*, *corridor*, *3s5z_vs_3s6z*). Only the *10m_vs_11m* scenario is homogeneous, where the army is composed of only a single unit type, while the others are heterogeneous.

The experimental results on different scenarios are shown in Figure 5. Compared with the state-of-the-art baseline methods, our proposed CIA successfully improves the final performance. In the easy homogeneous scenarios (*10m_vs_11m*), the strength gap between the agents and enemies is small. Thus, several baselines can also achieve promising results without complex tactics, while the benefit of diversity brought by CIA is not obvious. However, in the more difficult heterogeneous scenario (*2c_vs_64zg*, *7sz*), our method yields better performance than baselines. The results suggest that CIA can be utilized to provide a more discriminative credit assignment, which helps heterogeneous agents to explore diverse cooperation strategies. Moreover, this exploration may lead to a little drop in learning efficiency, but it is worthwhile as a trade-off for achieving non-trivial performance. To further test the potentiality of the proposed method, we also compare CIA in the super hard heterogeneous scenarios (*6h_vs_8z*, *corridor*, *3s5z_vs_3s6z*). In these challenging scenarios, learning distinguishable credits becomes complex due to the different unit types. Furthermore, there is a great disparity in strength between the two teams, and it is impossible to beat the enemies in a reckless way. The proposed CIA provides an impressive improvement in the performance over the baselines, showing its robustness to diversity promoting. Especially in the *3s5z_vs_3s6z* scenario, almost all compared baselines cannot learn any effective policy and perform poorly, while our method still maintains the superior performance.

Ablation Study

To verify the generalization of the proposed method, we apply the CIA module to two VD methods, including QMIX and QPLEX, and present the results in Figure 5. In general, the proposed CIA shows promising performance in enhancing different VD methods. Of special interest is the significant improvement in the super hard heterogeneous scenarios (*3s5z_vs_3s6z*), where credit assignment plays an important role. CIA (QMIX) and CIA (QPLEX) both obtain gratifying results superior to QMIX and QPLEX, indicating that the CIA module successfully enlarges the credit-level distinguishability. Moreover, it is notable that the improvement of incorporating CIA into different VD methods relies on the representation capability of the original models. The dedicated CIA module focuses on providing a regularization term to guide the optimization of the credit assignment with-




Figure 6: The performance comparison between the QMIX variants with the CC loss and the CIA loss, respectively.

out changing the intrinsic VD mechanism.

Moreover, to further demonstrate the advantage of CIA in promoting credit-level distinguishability, we also design a comparable variant that uses a credit classification loss (CC Loss) instead of the proposed contrastive identity-aware learning loss (CIA Loss). The credit classification loss is based on directly predicting the corresponding identity labels given the temporal credits. Figure 6 reports the experimental results of QMIX with different loss items. The dedicated CC loss can also make an improvement to alleviate the ambiguous credit assignment problem. In the simple scenarios (*10m_vs_11m*, *2c_vs_64zg*), the CIA loss offers close performance compared with CC loss. However, in the difficult scenarios (*corridor*, *3s5z_vs_3s6z*), the CIA loss outperforms the CC loss by a wide margin, while the CC loss suffers from a large variance. The results show that the CIA loss constraints the temporal credits and the identity representations into a specific credit-identity space without divergence, providing a more stable and robust optimization objective. However, the learned credits of the CC loss may distribute far away from the decision boundary constructed by the identity representations.

Conclusion

In this paper, we investigate the multi-agent diversity problem of the value decomposition (VD) paradigm in the MARL setting, observing that existing works on diverse agent networks are limited by the indistinguishable credit assignment of the VD network. Inspired by the observation, we propose a novel contrastive identity-aware learning (CIA) method to explicitly encourage the VD network to distinguish the assigned credits of multiple agents. Technically, we maximize the mutual information of the temporal credits and the identity representations of different agents with a tractable lower bound. We validate CIA over the StarCraft II micromanagement benchmark, and showcase that it yields results significantly superior to the state-of-the-art techniques in MARL. This simple yet effective CIA method further motivates us to explore the high-quality representation in future work rather than improving the network architecture. Another challenging direction is the theoretical study of the ambiguous credit assignment phenomenon. It is interesting to study the relationship between the IGM constraint and the ambiguous credit assignment problem. The IGM constraint may limit the distinguishable expressiveness of credit assignment in the existing VD architectures.

Acknowledgments

This work is funded by the National Key R&D Program of China (Grant No: 2018AAA0101503) and the Science and technology project of SGCC (State Grid Corporation of China): fundamental theory of human-in-the-loop hybrid-augmented intelligence for power grid dispatch and control.

References

- Ancona, M.; Ceolini, E.; Öztireli, C.; and Gross, M. 2018. Towards better understanding of gradient-based attribution methods for Deep Neural Networks. In *International Conference on Learning Representations*.
- Burda, Y.; Edwards, H.; Storkey, A. J.; and Klimov, O. 2019. Exploration by random network distillation. In *International Conference on Learning Representations*.
- Christianos, F.; Papoudakis, G.; Rahman, M. A.; and Albrecht, S. V. 2021. Scaling multi-agent reinforcement learning with selective parameter sharing. In *International Conference on Machine Learning*.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2019. Diversity is All You Need: Learning Skills without a Reward Function. In *International Conference on Learning Representations*.
- Fu, W.; Yu, C.; Xu, Z.; Yang, J.; and Wu, Y. 2022. Revisiting Some Common Practices in Cooperative Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*.
- Gu, P.; Zhao, M.; Hao, J.; and An, B. 2021. Online ad hoc teamwork under partial observability. In *International Conference on Learning Representations*.
- Hu, S.; Xie, C.; Liang, X.; and Chang, X. 2022. Policy Diagnosis via Measuring Role Diversity in Cooperative Multi-agent RL. In *International Conference on Machine Learning*.
- Jaderberg, M.; Czarnecki, W. M.; Dunning, I.; Marris, L.; et al. 2019. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443): 859–865.
- Jeon, J.; Kim, W.; Jung, W.; and Sung, Y. 2022. MASER: Multi-Agent Reinforcement Learning with Subgoals Generated from Experience Replay Buffer. In *International Conference on Machine Learning*.
- Jiang, J.; and Lu, Z. 2021. The Emergence of Individuality. In *International Conference on Machine Learning*.
- Lee, Y.; Yang, J.; and Lim, J. J. 2019. Learning to coordinate manipulation skills via skill behavior diversification. In *International conference on learning representations*.
- Li, C.; Wang, T.; Wu, C.; Zhao, Q.; Yang, J.; and Zhang, C. 2021. Celebrating Diversity in Shared Multi-Agent Reinforcement Learning. In *Annual Conference on Neural Information Processing Systems*.
- Long, Q.; Zhou, Z.; Gupta, A.; Fang, F.; Wu, Y.; and Wang, X. 2020. Evolutionary Population Curriculum for Scaling Multi-Agent Reinforcement Learning. In *International Conference on Learning Representations*.
- Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Annual Conference on Neural Information Processing Systems*.
- Luo, S.; Li, Y.; Li, J.; Kuang, K.; Liu, F.; Shao, Y.; and Wu, C. 2022. S2RL: Do We Really Need to Perceive All States in Deep Multi-Agent Reinforcement Learning? In *ACM Knowledge Discovery and Data Mining*.
- Macke, W.; Mirsky, R.; and Stone, P. 2021. Expected Value of Communication for Planning in Ad Hoc Teamwork. In *AAAI Conference on Artificial Intelligence*.
- Mahajan, A.; Rashid, T.; Samvelyan, M.; and Whiteson, S. 2019. MAVEN: Multi-Agent Variational Exploration. In *Annual Conference on Neural Information Processing Systems*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Rahman, M. A.; Hopner, N.; Christianos, F.; and Albrecht, S. V. 2021. Towards Open Ad Hoc Teamwork Using Graph-based Policy Learning. In *International Conference on Machine Learning*.
- Rashid, T.; Farquhar, G.; Peng, B.; and Whiteson, S. 2020. Weighted QMIX: Expanding Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *Annual Conference on Neural Information Processing Systems*.
- Rashid, T.; Samvelyan, M.; de Witt, C. S.; Farquhar, G.; Foerster, J. N.; and Whiteson, S. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*.
- Samvelyan, M.; Rashid, T.; de Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G. J.; Hung, C.; Torr, P. H. S.; Foerster, J. N.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. In *International Joint Conference on Autonomous Agents and Multi-agent Systems*.
- Son, K.; Kim, D.; Kang, W. J.; Hostallero, D.; and Yi, Y. 2019. QTRAN: Learning to Factorize with Transformation for Cooperative Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*.
- Sunehag, P.; Lever, G.; Grusly, A.; Czarnecki, W. M.; et al. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *International Joint Conference on Autonomous Agents and Multi-agent Systems*.
- Terry, J. K.; Grammel, N.; Hari, A.; Santos, L.; and Black, B. 2020. Revisiting parameter sharing in multi-agent deep reinforcement learning. *arXiv preprint arXiv:2005.13625*.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.

Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2021a. QPLEX: Duplex Dueling Multi-Agent Q-Learning. In *International Conference on Learning Representations*.

Wang, J.; Xu, W.; Gu, Y.; Song, W.; and Green, T. 2021b. Multi-Agent Reinforcement Learning for Active Voltage Control on Power Distribution Networks. In *Annual Conference on Neural Information Processing Systems*.

Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2021c. RODE: Learning Roles to Decompose Multi-Agent Tasks. In *International Conference on Learning Representations*.

Wang, T.; and Isola, P. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*.

Wang, T.; Wang, J.; Wu, Y.; and Zhang, C. 2020a. Influence-Based Multi-Agent Exploration. In *International Conference on Learning Representations*.

Wang, W.; Yang, T.; Liu, Y.; Hao, J.; Hao, X.; Hu, Y.; Chen, Y.; Fan, C.; and Gao, Y. 2020b. From Few to More: Large-Scale Dynamic Multiagent Curriculum Learning. In *AAAI Conference on Artificial Intelligence*.

Wu, T.; Zhou, P.; Liu, K.; Yuan, Y.; Wang, X.; Huang, H.; and Wu, D. O. 2020. Multi-agent deep reinforcement learning for urban traffic light control in vehicular networks. *IEEE Transactions on Vehicular Technology*, 69(8): 8243–8256.

Xu, X.; Jia, Y.; Xu, Y.; Xu, Z.; Chai, S.; and Lai, C. S. 2020. A multi-agent reinforcement learning-based data-driven method for home energy management. *IEEE Transactions on Smart Grid*, 11(4): 3201–3211.

Yan, Z.; and Xu, Y. 2020. A multi-agent deep reinforcement learning method for cooperative load frequency control of a multi-area power system. *IEEE Transactions on Power Systems*, 35(6): 4599–4608.

Yang, X.; Daquan, Z.; Liu, S.; Ye, J.; and Wang, X. 2022. Deep Model Reassembly. In *Annual Conference on Neural Information Processing Systems*.

Yang, X.; Ye, J.; and Wang, X. 2022. Factorizing knowledge in neural networks. In *European Conference on Computer Vision*.

Yu, C.; Velu, A.; Vinitsky, E.; Wang, Y.; Bayen, A.; and Wu, Y. 2021. The surprising effectiveness of ppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*.

Zhao, J.; Zhang, Y.; Hu, X.; Wang, W.; Zhou, W.; Hao, J.; Zhu, J.; and Li, H. 2022. Revisiting QMIX: Discriminative Credit Assignment by Gradient Entropy Regularization. *arXiv preprint arXiv:2202.04427*.

Zohar, R.; Mannor, S.; and Tennenholtz, G. 2022. Locality Matters: A Scalable Value Decomposition Approach for Cooperative Multi-Agent Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*.