

Adaptive Value Decomposition with Greedy Marginal Contribution Computation for Cooperative Multi-Agent Reinforcement Learning

Shanqi Liu
Zhejiang University
Hangzhou, China
shanqiliu@zju.edu.cn

Dong Xing
Zhejiang University
Hangzhou, China
dongxing@zju.edu.cn

Kun Kuang
Zhejiang University
Hangzhou, China
kunkuang@zju.edu.cn

Yujing Hu*
NetEase Fuxi AI Lab
Hangzhou, China
huyujing@corp.netease.com

Yu Xiong
NetEase Fuxi AI Lab
Hangzhou, China
xiongyu1@corp.netease.com

Yong Liu*
Zhejiang University
Hangzhou, China
yongliu@iipc.zju.edu.cn

Runze Wu
NetEase Fuxi AI Lab
Hangzhou, China
wurunze1@corp.netease.com

Changjie Fan
NetEase Fuxi AI Lab
Hangzhou, China
fanchangjie@corp.netease.com

ABSTRACT

Real-world cooperation often requires intensive coordination among agents simultaneously. This task has been extensively studied within the framework of cooperative multi-agent reinforcement learning (MARL), and value decomposition methods are among those cutting-edge solutions. However, traditional methods that learn the value function as a monotonic mixing of per-agent utilities cannot solve the tasks with non-monotonic returns. This hinders their application in generic scenarios. Recent methods tackle this problem from the perspective of implicit credit assignment by learning value functions with complete expressiveness or using additional structures to improve cooperation. However, they are either difficult to learn due to large joint action spaces or insufficient to capture the complicated interactions among agents which are essential to solving tasks with non-monotonic returns. Moreover, applications in real-world scenarios usually require policies to be interpretable, but interpretability is limited in the implicit credit assignment methods. To address these problems, we propose a novel explicit credit assignment method to address the non-monotonic problem. Our method, Adaptive Value decomposition with Greedy Marginal contribution (AVGM), is based on an adaptive value decomposition that learns the cooperative value of a group of dynamically changing agents. We first illustrate that the proposed value decomposition can consider the complicated interactions among agents and is feasible to learn in large-scale scenarios. Then, our method uses a greedy marginal contribution computed from the value decomposition as an individual credit to incentivize agents to learn the optimal cooperative policy. We further extend the module with an action encoder to guarantee the linear time complexity for computing the greedy marginal contribution. Experimental results demonstrate that our method achieves significant performance improvements

in several non-monotonic domains. Besides, we showcase that our model maintains a good sense of interpretability and rationality. This suggests our model can be applied to scenarios with more realistic demands.

KEYWORDS

Multi-Agent Cooperation; Credit Assignment; Non-Monotonic

ACM Reference Format:

Shanqi Liu, Yujing Hu*, Runze Wu, Dong Xing, Yu Xiong, Changjie Fan, Kun Kuang, and Yong Liu*. 2023. Adaptive Value Decomposition with Greedy Marginal Contribution Computation for Cooperative Multi-Agent Reinforcement Learning. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 12 pages.

1 INTRODUCTION

Many real-world tasks like bimanual manipulation [11], autonomous driving [13] and swarms [8] require the cooperation of multiple agents. Especially, in many scenarios agents are expected to choose the optimal actions simultaneously to complete the common goal, such as the bimanual lifting task requiring dual arms to lift the object simultaneously [1, 2]. Learning cooperative policies in these tasks remains challenging as agents' joint action spaces grow exponentially with the number of agents and performing reward decomposition, known as credit assignment, is challenging due to complex interactions among agents. Cooperative multi-agent reinforcement learning (MARL) has been broadly used to learn effective behaviors in such tasks from agents' experiences. Currently, a popular paradigm for value-based cooperative MARL is centralized training with decentralized execution (CTDE), and representative methods include VDN and QMIX [20, 23]. These methods learn the centralized value functions as monotonic factorizations of each

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

*The corresponding authors.

agent’s utility function and enable decentralized execution by maximizing each agent’s corresponding utility function, known as the Individual Global Maximum (IGM) principle [6]. Despite the fact that these methods have successfully solved tasks with monotonic payoff matrices optimally [7], they fail in tasks with non-monotonic payoff matrices. As the monotonic value function restricts the value function to sub-optimal value approximations in environments with non-monotonic payoffs [22, 25], they cannot represent the policy that an agent’s optimal action depends on actions from other agents. This problem, known as the relative overgeneralization [17], prevents the agents from solving tasks such as bimanual lifting.

To address this problem, recent works either aim to learn value functions with complete expressiveness capacity [25] or use auxiliary approaches such as placing more importance on better joint actions to find the optimal cooperative policy [19]. However, they still face other problems. On the one hand, learning complete expressiveness is difficult since the joint action spaces are exponentially related to the number of agents. On the other hand, using auxiliary approaches to improve cooperation cannot capture the complicated interactions among agents and lacks theoretical optimal guarantees [5] (Section 3 details the reasons for the problems). Furthermore, current methods are usually implicit credit assignment methods, which learn the credit assignment by using a mixing network. Compared to the explicit credit assignment methods based on Shapley Value [21] or marginal contribution, implicit methods lack interpretability for the distributed credit, which hinders their application in scenarios where security is an essential property.

In this work, we propose a novel explicit credit assignment method that solves the non-monotonic problem by using the greedy marginal contribution of each agent. First, we propose an adaptive value decomposition method that models individual utility as a distribution over actions of other observable agents. The value decomposition enables agents to learn the individual utility that can take into account the interactions with others. We illustrate that such an individual utility can represent non-monotonic payoff matrices correctly and we further investigate why previous methods that model the utility as a specific value fail in the non-monotonic settings in Section 3. Furthermore, since the utility merely models interactions among observable agents, the learning complexity does not increase exponentially with the number of agents as in methods that learn complete expressiveness. However, although such individual utilities can describe the payoff matrices correctly, we cannot directly use it as policy’s value function as it requires other agents’ actions. To address this problem, our method optimizes each agent’s decentralized policy by learning a greedy marginal contribution from the adaptive value decomposition, which calculates the optimal marginal contribution to avoid the requirement to know the other agents’ actions for policy computation. We illustrate that optimizing the greedy marginal contribution can achieve the cooperative policy by avoiding taking actions that lead to mis-coordination. This also enhances exploration as the number of samples taking cooperative actions increases. Furthermore, since the computation time complexity of maximizing the marginal contribution is exponentially related to the number of observable agents, we propose an action encoder that maps actions into a condensed latent space to guarantee a linear searching time for the maximization. This reduces computational burden and makes the

learning process practical. Finally, our method also increases the adaptability of learned policy to cooperate with scalable agents as our method learns the values of actions cooperating with dynamically changing observable agents. We evaluate our method against several state-of-the-art baselines in non-monotonic robotics tasks that require simultaneous coordination amongst agents. The result shows that our method achieves significant improvements in asymptotic performance and efficiency. Moreover, we showcase the interpretability of our method that current methods lack.

2 RELATED WORK

Many general deep MARL methods have been used in complex multi-agent environments, including COMA [3], MADDPG [14], etc. Currently, CTDE forms the de facto mainstream paradigm in cooperative MARL [9, 14]. Regarding the details, VDN [23] learns the joint-action Q-values by factoring them as the sum of each agent’s utility. QMIX [20] extends VDN to allow the joint action Q-value to be a monotonic combination of each agent’s utility that can vary depending on the global state. However, the monotonic constraints on the joint action-values introduced by QMIX and related QMIX-based variant methods result in provably poor exploration and relative overgeneralization [17]. To address this problem, QPLEX [25] and QTRAN [22] aim to learn value functions with complete expressiveness capacity. However, reports are that they perform poorly when being used in practice [5, 24]. This is because learning the complete expressiveness is impractical in complicated MARL tasks due to the challenging exploration in large joint action spaces. In contrast, we restrict the modeling of agent interaction value to be within pairs of potentially cooperating agents. This prevents the complexity of learning from growing exponentially along with the size of joint action spaces. Moreover, our method also enhances exploration by increasing the probability of sampling cooperative actions. There also exist methods using auxiliary structures to solve the problem. For instance, MAVEN [15] hybridises value and policy-based methods by introducing a latent space for hierarchical control. This allows MAVEN to achieve committed, temporally extended exploration. Weighted QMIX [19] is based on QMIX and rectifies the suboptimality by introducing weights to place more importance on the better joint actions. UneVEN [5] learns a set of related tasks simultaneously with a linear decomposition of universal successor features. However, these methods cannot estimate the value of actions considering the changes in other agents’ actions and usually fail in environments with extreme rewards. On the contrary, our method can estimate the contribution of each agent according to the utility considering the interactions of all cooperative agents which is shown to work in all non-monotonic environments. Moreover, all the methods mentioned above belong to implicit credit assignment methods, which lack interpretability for the distributed credits.

Our work is also related to credit assignment. Compared to implicit credit assignment methods, explicit methods attribute each agent’s contribution to being at least provably locally optimal [10]. COMA [3] utilizes a counterfactual advantage to learn the value function. QPD [28] designs a multi-channel mixer critic and leverages integrated gradients to distribute credits along paths. However, they use monotonic value functions and ignore the complex




Figure 1: Left: The illustration of Eq. (17). The orange area is the payoff matrixes that can be solved by methods using $Q_i(\tau_i, a_i)$ as policy. Right: Payoff matrix of the bimanual lifting example.

interactions between agents, which are critical for cooperation in non-monotonic tasks. To address this, SQDDPG [26] and Shapley [12] use Shapley Value [21] to estimate the complex interactions between agents. Shapley Value originates from cooperative game theory and is able to distribute benefits reasonably by estimating the contribution of participating agents. However, SQDDPG [26] assumes that the agents take actions sequentially, which does not hold in many scenarios. Shapley [12] uses a single centralized critic network to estimate Shapley Value, which is infeasible when the number of agents is large. Moreover, these methods can only get approximated Shapley Value as calculating the Shapley Value involves exponential time complexity [26].

3 MOTIVATION

In this section, we use a bimanual task in robotics as a motivating example, although our method can be applied to general tasks with non-monotonic payoffs. The bimanual tasks are defined as tasks requiring a dual-arm robot to manipulate. Compared to a single-arm robot, a dual-arm robot has prominent advantages in finishing complex real-world tasks, especially humanly intuitive tasks [18, 27]. For example, simple tasks such as lifting a single heavy object or complex tasks such as outer-space assembly and repair or domestic work [11]. However, learning the bimanual tasks using MARL is challenging since the cooperative actions between multiple arms usually depend on each other’s actions. Furthermore, the payoff matrixes of these tasks are non-monotonic which prevent agents from learning the optimal policy. To illustrate the reason for such a problem, we use the bimanual lifting task as an example. The task requires two agents to take the cooperative action C simultaneously to lift the object, and the collaborative reward is $+R$. Otherwise, the agent will receive a punishment $-P$ for taking action C alone since the object cannot be lifted alone and the energy is wasted. Meanwhile, action L represents lazy actions (e.g., staying still) in the action spaces that cannot achieve cooperation but avoid punishment. The corresponding non-monotonic payoff matrix is in Figure 1.

We investigate the reason for learning difficulties in such tasks by analyzing the individual utility form. To learn the decentralized policy, current methods model each agent’s individual utility as a certain value function $Q_i(\tau_i, a_i)$. However, since the returns depend on the interactions between all agents in such tasks, such a decomposition cannot consider others’ actions and is insufficient

to represent the optimal policy in some situations. We illustrate that a more comprehensive individual utility that can consider the interactions between agents is $Q_c^i(\tau_i, u_i^-, a_i)$ where u_i^- is joint actions of potential cooperative agents (discussed in detail in Section 5.1). According to this decomposition, the individual utility $Q_i(\tau_i, a_i)$ should be viewed as a variable sampled from the distribution $Q_c^i(\tau_i, u_i^-, a_i)$ over u_i^- . Optimization of $Q_i(\tau_i, a_i)$ instead of $Q_c^i(\tau_i, u_i^-, a_i)$ can cause the policy converge to an average policy. We illustrate that this explains why current methods fail in environments with non-monotonic payoff matrices such as bimanual lifting. For instance, in the bimanual lifting task, the learned policy will not be able to represent the optimal policy when

$$\frac{R}{P} < \frac{2p_l - 1}{1 - p_l}. \quad (1)$$

where p_l is the probability of each policy taking action L . The detailed derivation is included in Appendix 6. Figure 1 shows the result of Eq. (17).

We find that methods directly using $Q_i(\tau_i, a_i)$ as policy such as QMIX can only solve a small part of all possible payoffs since even a small P can cause the policy to fail because the weight of P grows exponentially as p_l grows. However, other methods proposed to tackle the relative overgeneralization can decrease the p_l to solve more payoffs. For instance, Weighted QMIX decreases p_l by placing more importance on the better joint actions. However, since it does not model the actual interactions among agents, the importance weight is empirical. Therefore, it cannot promise the p_l to decrease to zero and still fails when P is large. QPLEX can reduce p_l to zero by learning converged value functions with complete expressiveness capacity. However, such value functions are hard to learn as joint action spaces grow exponentially as the number of agents grows. Moreover, learning complete expressiveness requires sufficient exploration which is difficult to achieve in the example environment, since the p_l is usually large during uniform exploration. This is because there are only a few cooperative actions (e.g., lift) and the lazy actions (e.g., stay and move) are major in the entire action spaces. However, our method tackles these problems by learning a feasible value decomposition which can consider the interactions among agents and using a greedy marginal contribution to guarantee that the p_l can be reduced to zero. Therefore, our method achieves both optimality and efficiency in practice.

4 PRELIMINARY

4.1 Dec-POMDP

A fully cooperative multi-agent sequential decision-making task can be described as a decentralized partially observable Markov decision process (Dec-POMDP), which is defined by a set of possible global states S , actions A_1, \dots, A_N , and observations $\Omega_1, \dots, \Omega_N$. At each time step, each agent $i \in \{1, \dots, N\}$ chooses an action $a_i \in A_i$, and they together form a joint action $\mathbf{u} \in U$. The next state is determined by a transition function $P : S \times U \rightarrow S$. The next observation of each agent $o_i \in \Omega_i$ is updated by an observation function $O : S \rightarrow \Omega$. All agents share the same reward $r : S \times U \rightarrow \mathbb{R}$ and a joint value function $Q_{tot} = E_{S_{t+1}, \dots, a_{t+1}, \dots} [R_t | s_t, \mathbf{u}_t]$ where $R_t = \sum_{j=0}^{\infty} \gamma^j r_{t+j}$ is the discounted return. The observation of each agent can also be replaced by the history of actions and observations

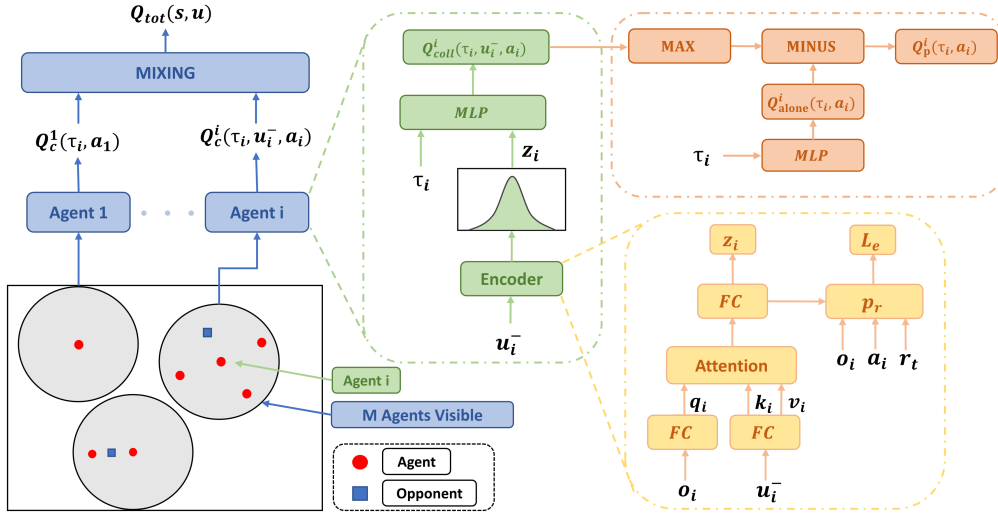


Figure 2: The architecture of our method. Left (Blue): The architecture for total Q-values. Middle (Green): The framework of the non-monotonic critic $Q_c^i(\tau_i, u_i^-, a_i)$. Upper Right (Red): The framework of greedy marginal contribution and $Q_p^i(\tau_i, a_i)$. Lower Right (Yellow): The framework of action encoder and predictive model.

of each agent as a proxy to handle partial observability [20, 23]. The history of actions and observations of agent i can be viewed as τ_i which is $(o_i^0, a_i^0, \dots, o_i^t)$.

4.2 Marginal Contribution

In this work, we also introduce the marginal contribution of Shapley Value. The marginal contribution in Shapley Value of agent i is defined as

$$\phi_i = v(C) - v(C/i) \quad (2)$$

where C is a team consisting of agents cooperating with each other to achieve a common goal and C/i represents the set with the absence of agent i . $v(C)$ refers to the value function for estimating the cooperation of a set of agents. Additionally, we use N, M_i to refer to the total number of agents and M_i is the number of agents observed by agent i .

5 METHOD

In this section, we propose an explicit credit assignment method, Adaptive Value decomposition with Greedy Marginal contribution (AVGM), to learn the optimal cooperative policy in tasks with non-monotonic payoffs. First, we propose a value decomposition which can capture the interactions between agents. Then, we propose an actor-critic structure based on the value decomposition that trains a centralized critic and calculates a greedy marginal contribution from the critic to train a decentralized execution policy.

5.1 Learning Non-Monotonic Critic via Adaptive Value Decomposition

Since current individual utility $Q_i(\tau_i, a_i)$ fails to capture the interactions between agents, we revise the agent's individual utility to a more comprehensive form. First of all, as the total rewards of environment are generated by all kinds of possible interactions

between agents, we have

$$r_{tot}(s, u) = \sum_{i=1}^N \hat{r}_1(o_i, a_i) + \sum_{\substack{i < j = N \\ i < j = 2}} \hat{r}_2(o_i, o_j, a_i, a_j) + \dots + \hat{r}_N(o_1, \dots, o_N, a_1, \dots, a_N). \quad (3)$$

Notable, each item \hat{r}_m in Eq. (14) exists only when there are interactions among the referred agents. Otherwise, it should be neglected as the reward is represented by other items. However, the enormous number of combinations makes it difficult to use this decomposition in practice. Fortunately, we can simplify the decomposition in multi-agent decentralized execution settings. In decentralized execution settings, we assume the cooperative mode is that each agent should only interact with others who are observable to the agent. The reason for this is that the decentralized policy can only take action according to information within the observation. If an agent needs to cooperate with other agents that are not within its observation field, we must introduce a communication method to transfer the necessary messages. Otherwise, the agents cannot identify which agents can potentially achieve cooperation, which would make the examples of cooperation irregular from the view of local observations and makes it impossible to learn the cooperative policy. Additionally, most robotics scenarios have a smaller range of interaction than the field of view in practice, which means that if an agent is within the interaction range of another agent, it must also be within its field of view.

Following this principle, we define team i as all M_i agents in the view field of agent i . The total reward of agent i can be described as $r_{team}^i(o_i, u_i^-, a_i)$, where u_i^- means the joint actions of the rest of agents in team i . Since observation o_i can serve as the global state of the team which has included all other agents' information and agent i only requires others' actions information to refer to the total reward. Furthermore, considering all kinds of possible interactions

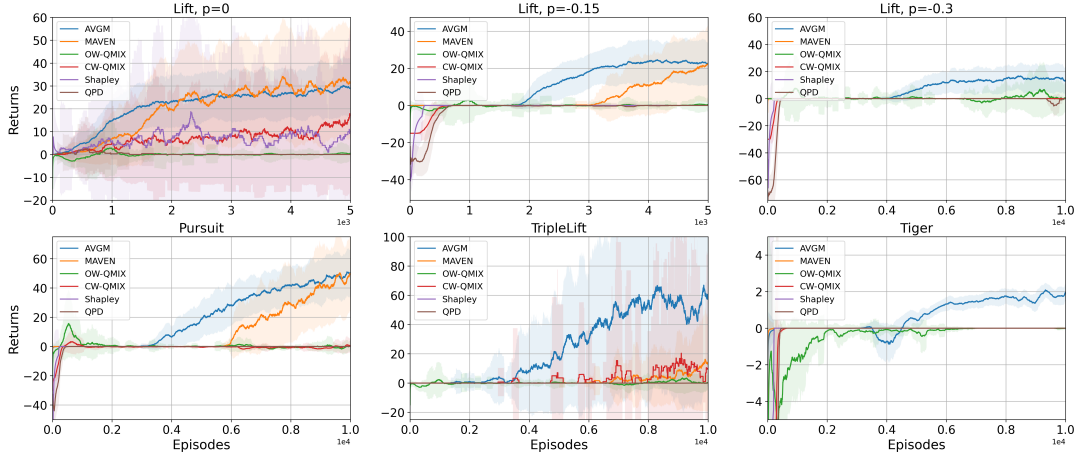


Figure 3: Top: The results in *lift* with the penalty growing from 0 to -0.3. Bottom: The results in *pursuit*, *triplelift* and *tiger*.

between agents, we can decompose the $r_{team}^i(o_i, u_i^-, a_i)$ as Eq. (14),

$$r_{team}^i(o_i, u_i^-, a_i) = \hat{r}_1^i(o_i, a_i) + \sum_{j=1, j \neq i}^{M_i} \hat{r}_2^i(o_i, a_j, a_i) + \dots + \hat{r}_{M_i}^i(o_i, a_1, \dots, a_{M_i}, a_i). \quad (4)$$

where \hat{r}_m^i exists only when there are interactions between the described agents and agent i , otherwise it should be equal to zero. Combining Eq. (15) with Eq. (14), we have

$$r_{tot}(s, u) = \sum_{i=1}^N r_{team}^i(o_i, u_i^-, a_i). \quad (5)$$

The detailed derivation can be found in Appendix 5. Based on this decomposition, we can derive the corresponding adaptive value decomposition, using $Q_c^i(\tau_i, u_i^-, a_i)$ as each agent's individual utility which learns the value of cooperating with the dynamically changing observable agents at each time step. Then, we have

Theorem 1: For any $r_{tot}(s, u)$, the corresponding $Q_{tot}(s, u) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r_{tot}(s, u) | \pi]$ and each agent's utility $Q_c^i(\tau_i, u_i^-, a_i)$ satisfies

$$\arg \max_u (Q_{tot}(s, u)) = \{ \arg \max_{a_1} (Q_c^1(\tau_1, u_1^-, a_1)), \dots, \arg \max_{a_N} (Q_c^N(\tau_N, u_N^-, a_N)) \}. \quad (6)$$

Detailed proof can be found in Appendix 5. Theorem 1 indicates that the adaptive value decomposition using utility $Q_c^i(\tau_i, u_i^-, a_i)$ can represent the unbiased value decomposition given any reward function and satisfies the IGM principle for decentralized execution. Furthermore, since the utility only involves observable agents, it does not make the learning problem harder along with the increasing of agent number within the environment.

Since we have proposed $Q_c^i(\tau_i, u_i^-, a_i; \theta)$ as an agent's utility to overcome the non-monotonic problem. We would like to learn a policy from it that can cooperate considering each other's actions. However, the policy $a_i \sim Q_c^i(\tau_i, u_i^-)$ requires u_i^- to be taken but u_i^- also requires a_i to be produced, which is a deadlock situation. To address this problem, we propose an actor-critic structure that trains

a centralized critic and calculates a greedy marginal contribution from the critic to train a decentralized execution policy.

Following Theorem. 1, we construct our non-monotonic centralized critic. Since Theorem 1 indicates that Q_{coll}^i can represent the unbiased value decomposition given any reward function and satisfies the IGM principle for decentralized execution. We can use all Q_{coll}^i to calculate Q_{tot} through a monotonic mixing network similar to QMIX, and Theorem 1 promises the mixing value is unbiased. Therefore, our centralized critic consists of each agent's adaptive utility $Q_c^i(\tau_i, u_i^-, a_i)$ and a mixing network to produce the global Q-value $Q_{tot}(s, u)$. The optimization objective is mean squared error (MSE)

$$\mathcal{L}_{TD}(\theta) = \mathbb{E}_{\pi} [Q_{tot}(s_t, \mathbf{u}_t) - y_t]^2 \quad (7)$$

$$y_t = r_t + \gamma \max_{\mathbf{u}_{t+1}} Q_{tot}(s_{t+1}, \mathbf{u}_{t+1}).$$

However, we notice that agents are not always able to observe other agents and the value function when agents have no interactions with each other should be different. Therefore, we construct $Q_c^i(\tau_i, u_i^-, a_i)$ as two value functions,

$$Q_c^i(\tau_i, u_i^-, a_i) = \begin{cases} Q_{coll}^i(\tau_i, u_i^-, a_i; \theta_1) & M_i > 0 \\ Q_{alone}^i(\tau_i, a_i; \theta_2) & M_i = 0 \end{cases} \quad (8)$$

where M_i is the number of agents observed by agent i . $Q_{coll}^i(\tau_i, u_i^-, a_i)$ represents the cooperative value function and is selected when any other agent can be observed. The $Q_{alone}^i(\tau_i, a_i)$ learns the selfish value function of an agent acting alone and is selected when no agents are around. During training process, we choose one of these two value functions as $Q_c^i(\tau_i, u_i^-, a_i)$ according to whether any other agent can be observed at each time step to produce $Q_{tot}(s, u)$.

5.2 Explicit Credit Assignment via Greedy Marginal Contribution

Although we already have a non-monotonic critic function, we still need to learn a decentralized execution policy. Fortunately, we can train a decentralized policy by optimizing the marginal contribution of each agent to avoid requiring others' actions. The




Figure 4: Top: The results of the action encoder outputs from initial (left) to final (right) of the training process. The x-axis and y-axis are the actions of other agents. Different colors and numbers represent different action representations. Bottom: The results of credit assignment values from initial (left) to final (right) of the training process. All results have been normalized for better representation.

reason is that the marginal contribution can only be maximized by achieving cooperation. Therefore, if an action can maximize the marginal contribution, it is the cooperative action. According to the definition of marginal contribution, if agent i can observe other agents, we have

$$\phi_i(\tau_i, a_i) = v(T_i) - v(T_i/i) = Q_{coll}^i(\tau_i, u_i^-, a_i) - Q_{coll}^i(\tau_i, u_i^-, a_i^l) \quad (9)$$

where T_i is the set of agents consisting of agent i and all agents observed by agent i , a_i^l is the lazy action which cannot lead to cooperation with any agent. In practice, we use

$$a_i^l = \arg \max_{a_i} (Q_{alone}^i(\tau_i, a_i)) \quad (10)$$

The rationale is that the policy represented by Q_{alone} is a selfish policy only considering acting alone. In non-monotonic environments, such a policy is supposed to take lazy actions to avoid the punishment in payoff matrix. Moreover, even in monotonic environments, the cooperative actions are usually different from selfish actions, so Eq. 23 holds in all kinds of environments. However, there is still another problem that for any (τ_i, a_i) , there are multiple $\phi_i(\tau_i, a_i)$ corresponding to different u_i^- . This can cause our policy to converge to an average policy like the situation we discussed in Section 3. To address this issue, we propose a greedy marginal contribution,

$$\begin{aligned} \phi_i^*(\tau_i, a_i) &= Q_{coll}^i(\tau_i, u_i^{*-}, a_i) - Q_{coll}^i(\tau_i, u_i^{*-}, a_i^l) \\ (u_i^{*-}, a_i^*) &= \arg \max_{(u_i^-, a_i)} (Q_{coll}^i(\tau_i, u_i^-, a_i)) \end{aligned} \quad (11)$$

The insight is that we expect our policy to learn the potential optimal value of action a_i . The optimal value can be reached when all other agents take the optimal cooperative actions u_i^{*-} to cooperate with agent i . In other words, we encourage the agent to learn each action's value based on the optimistic belief that the possibility of other agents to take non-cooperative action is zero so that all other agents would cooperate with itself. As the main problem in

the non-monotonic environment is that agents tend to take lazy actions fearing other agents not cooperating, this optimistic belief can facilitate exploration by increasing the probability of sampling cooperative actions and promises that agents will jump off the sub-optimal policy to converge to the optimal cooperative policy. In this way, we have our decentralized policy's value function $Q_p^i(\tau_i, a_i; \mu)$ and decentralized policy as $\pi(a_i|\tau_i) = \arg \max_{a_i} (Q_p^i(\tau_i, a_i))$. We use $\phi_i^*(\tau_i, a_i)$ as target when an agent can observe other agents and use Q_{alone} when acting alone. The overall loss objective is

$$\mathcal{L}_p(\mu) = \begin{cases} \mathbb{E}_{\pi} [Q_p^i(\tau_i, a_i) - \phi_i^*(\tau_i, a_i)]^2 & M_i > 0 \\ \mathbb{E}_{\pi} [Q_p^i(\tau_i, a_i) - Q_{alone}(\tau_i, a_i)]^2 & M_i = 0 \end{cases} \quad (12)$$

5.3 Learning Actions Category Representation

However, the structure of $Q_c^i(\tau_i, u_i^-, a_i)$ is infeasible to use in practice because u_i^- 's dimension is uncertain that we cannot construct a neural network for it. More importantly, we need to search for the maximization of $Q_c^i(\tau_i, u_i^-, a_i)$ when calculating our greedy marginal contribution. The time complexity will be exponential if we search all possible u_i^- . To address these problems, we propose an action encoder $z_i \sim f(u_i^-; \phi)$ to map the actions' information into different categories in latent space to reduce the dimension. The insight behind this is that we find many of the combinations in u_i^- actually have the same effect on the environment and agent i , so they should have the same category in the latent space. For example, in a scenario requiring four agents to take cooperative actions, agent i would face the same situation regardless of which of the other three agents takes the lazy action.

To handle the uncertain dimension of u_i^- , we use an attention structure [4, 16] in the encoder. Moreover, the attention structure can help agents pay more attention to agents that are more likely to cooperate. We calculate the weight of each a_j in u_i^- according to the information of o_i . The overall structure of attention is shown in Figure 2. However, since we lack labels for the categories of




Figure 5: Left: Comparison between different z dimensions. Right: Comparison between different agents’ choosing methods.

the actions, the encoder has to learn in an unsupervised manner. Therefore, in practice we use a predictive model $p_r(o_i, f(u_i^-), a_i; \eta)$ to learn the encoder, as the learned action categories should contain enough information such that the next reward of agent i can be predicted when given the observations and actions of agent i . The overall model is trained by minimizing the following loss function which is derived from Eq. (16),

$$\mathcal{L}_e(\phi, \eta) = \left(\sum_{i=0}^N p_r(o_i, f(u_i^-), a_i) - r_t \right)^2 \quad (13)$$

where r_t is the global reward. At last, we modify $Q_c^i(\tau_i, u_i^-, a_i)$ by replacing u_i^- with $z_i \sim f(u_i^-)$ to $Q_c^i(\tau_i, z_i, a_i)$. This formula promises a linear search time $O(Z)$ where Z is the pre-defined number of categories.

6 EXPERIMENT

6.1 Environments and Experimental Settings

We design multiple challenging cooperative robotics tasks with non-monotonic payoffs to evaluate each method. The experiments are conducted based on MAgent [29], where we implement four tasks with different focuses: *lift*, *triplelift*, *pursuit* and *tiger*. All tasks are non-monotonic environments which the tasks require all agents to take optimal cooperative actions based on others’ actions. Otherwise, agents will receive penalty. *lift* and *triplelift* require two or more and three or more agents to lift a cargo together which are similar to the bimanual lifting task but more tricky because agents need to move to the cargo first. *pursuit* is a predator-prey task that needs two or more agents to attack a prey together to capture it. *tiger* is more difficult than *pursuit* as its prey will be killed by agent’s attack actions and the prey in *tiger* can recover health at each time step, which encourages agents to learn to attack prey with higher health and refrain from killing prey to get more rewards. We use the performance of evaluation episodes with greedy action selections as the final performance. The performance is evaluated by the average return on each agent. All experiments are carried out with five random seeds. More details about the implementations of all scenarios are included in Appendix 1.

In the experiments, we compare our method AVGM with QPD, Shapley, MAVEN and WQMIX. All methods use the same basic hyperparameters and network structures with similar parameters to ensure the comparison is fair. Please refer to Appendix 2 for more experimental settings and <https://github.com/Locke637/AVGM> for the codes.




Figure 6: The results in large-scale scenarios. Left: Results in *pursuit*. Right: Results in *lift*.

6.2 Performance

We show the results of six different robotics scenarios in Figure 3. In *lift*, we demonstrate how penalties affect the learning process of all methods. We evaluate all methods in *lift* with the penalty growing from 0 to -0.3. The results indicate that only AVGM can solve all scenarios while other methods’ performance decreases as the penalty grows. Especially, the results show that other methods tend to learn a lazy policy that never takes cooperative actions to avoid the penalty when penalty grows larger. This phenomenon is consistent with our previous analysis in Section 3. We also notice that MAVEN and WQMIX have better performance in scenarios with larger penalties as they are designed to alleviate non-monotonic. However, they cannot solve the non-monotonic problem completely as they fail in tasks with the most significant penalty. Furthermore, the poor performance of QPD is due to the lack of considering interactions between agents, while Shapley fails in harder scenarios because the centralized critic is not sufficient to represent the true marginal contribution of optimal action, thus leading to a sub-optimal policy. In *pursuit*, the result is similar as in *lift* that only AVGM learns an optimal policy. In *triplelift*, we find that the performance of other methods diminishes as cooperation becomes harder to achieve. However, AVGM still solves the task by considering others’ actions and searching for the optimal marginal contribution. Finally, all other methods fail in *tiger* as inefficient exploration would kill all prey in this task, making it difficult to find a cooperative mode. The result indicates that AVGM is highly efficient in learning in non-monotonic environments.

6.3 Analysis of Interpretability

In this section, we analyze several properties of our method to show the interpretability of AVGM. To better demonstrate, we choose a typical situation in *lift* with penalty -0.3, where all agents are in the right position and ready to lift the cargo. The schematic diagram is included in Appendix 3. This situation is suitable for demonstrating interpretability as agents must consider each other’s actions to decide whether to take cooperative or lazy actions.

In the top of Figure 4, we demonstrate the output changes of the action encoder during training. We illustrate that the action encoder learns a meaningful category representation of others’ actions. The representation classifies others’ actions according to whether they take lift actions or movement actions. We notice that the encoder first learns the easy situations that both agents are lifting or moving. Through further training, the encoder learns to classify the rest combinations of actions.




Figure 7: Results of scalability. Left: Results in *pursuit*. Right: Results in *tiger*.

Furthermore, we show the interpretability of our marginal contribution at the bottom. We compare the marginal contribution of taking optimal lift action and lazy stay action under the condition that others are taking cooperative actions or lazy actions. The result shows that the marginal contribution of taking lift action has an increasingly large advantage given other’s lift actions over given other’s lazy actions as the training progresses. Meanwhile, the marginal contribution of taking stay action when others are taking lazy actions is larger than that of others taking cooperative actions. The results show that the marginal contribution is reasonable as it is consistent with the analysis using prior knowledge. Therefore, we showcase the interpretability of our method which guarantees further potential applicants in the real-world scenarios since we can explain the behavior of each agent and reduce the unpredictable risks. Finally, we also demonstrate the variance of all agents’ $Q_c^i(\tau_i, u_i^*, a_i^*)$ during training, which shows that the agents would learn a consistent perception of the environment when considering others’ actions in Appendix 4.

6.4 Ablations

6.4.1 Comparison among Using Different z Dimensions. We compare the performance of AVGM using different z dimensions (8, 16, 24, 32, 40) in *pursuit*. The results are shown in Figure 5. The results indicate that using different z dimensions has a limited impact on the final performance when the number of dimensions is over 16. However, using 8-dimension z can decrease the performance significantly as such low-dimensional latent action embedding is insufficient to represent all the necessary behavior modes. This showcases the robustness of AVGM over the hyperparameter of z dimensions within a reasonable range.

6.4.2 Comparison among Using Different Agents’ Choosing Methods. We compare the performance of using all agents, using only visible agents and using part of visible agents in *pursuit*. The results are shown in Figure 5. Specifically, using part of visible agents means that we randomly abandon half of the visible agents. The results show that using all agents and only visible agents has similar performance. This indicates that the information among visible agents is enough to represent the optimal policy, which is in line with our analysis. Meanwhile, the performance of using part of visible agents decreases significantly. This indicates that the absence of the joint actions of visible agents u_i^- leads to a suboptimal policy.

6.4.3 Scalability. First, we compare the performance in large-scale scenarios with more agents (4, 8, 12) in *pursuit* and (3, 8, 12) in *lift*.

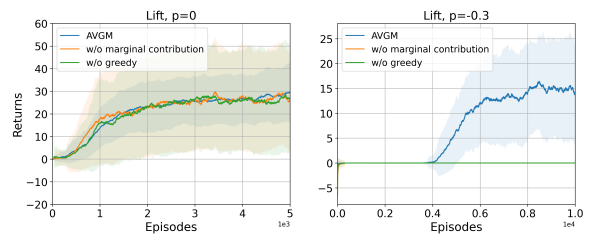


Figure 8: Results of ablation studies for marginal contribution. Left: Results in *lift* with penalty of 0. Right: Results in *lift* with penalty of -0.3.

The results are shown in Figure 6. The results show that the final performance is similar in all settings. This indicates that adding agents has limited influence on the performance of AVGM, as the number of agents in the view of each agent is much smaller and limited compared with the total number of all agents in the environment. This demonstrates that AVGM can learn in large-scale scenarios that are intractable for methods that learn value functions with complete expressiveness. Moreover, training with more agents can learn cooperation at the early stages since the policy is trained with more experience and the odds of discovering the cooperative mode become higher. Additionally, since the structure of AVGM can handle variable input dimensions and the policy can learn to cooperate with scalable agents, we investigate whether the learned policy can be zero-shot transferred across scenarios with different numbers of agents. We test the scalability of our method in *pursuit* and *tiger* with a varying agents’ number of 3-5. The results in Figure 7 indicate that our method can generalize to scenarios with different numbers of agents because there is no significant performance degradation in the test environments.

6.4.4 Ablations for Marginal Contribution. We also conduct ablation studies on the marginal contribution, including using the actual marginal contribution instead of the greedy one and not using marginal contribution by optimizing policy directly through using $Q_c^i(\tau_i, u_i^*, a_i)$. The ablation studies are conducted in *lift* with penalties of 0 and -0.3. The results are shown in Figure 8. We find that lacking these parts of our method leads to significant reductions in non-monotonic scenarios, while the influence is limited in monotonic ones.

7 CONCLUSION

In this work, we propose a novel explicit credit assignment method to solve the robotics tasks with non-monotonic returns. Our method, AVGM, is based on an adaptive value decomposition considering other observable agents’ actions and uses greedy marginal contribution to encourage agents to learn the optimal cooperative policy. This enables our method to learn cooperation in tasks with non-monotonic returns, which is a challenge for existing value-based CTDE methods. Moreover, we propose an action encoder to guarantee that the time complexity of AVGM is linear. The experiments show significant performance improvement and the interpretability of AVGM.

8 ACKNOWLEDGMENTS

The publication was supported by NSFC 62088101 Autonomous Intelligent Unmanned Systems and by a Grant from The National Natural Science Foundation of China (No. U21A20484).

REFERENCES

- [1] RC Bonitz and Tien C Hsia. 1996. Internal force-based impedance control for cooperating manipulators. *IEEE Transactions on Robotics and Automation* 12, 1 (1996), 78–89.
- [2] Fabrizio Caccavale, Pasquale Chiacchio, Alessandro Marino, and Luigi Villani. 2008. Six-dof impedance control of dual-arm cooperative manipulators. *IEEE/ASME Transactions On Mechatronics* 13, 5 (2008), 576–586.
- [3] 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.
- [4] Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural Turing machines. *arXiv preprint arXiv:1410.5401* (2014).
- [5] Tarun Gupta, Anuj Mahajan, Bei Peng, Wendelin Böhmer, and Shimon Whiteson. 2021. Uneven: Universal value exploration for multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 3930–3941.
- [6] Wan Ju Kang David Earl Hostallero, Kyunghwan Son, Daewoo Kim, and Yung Yi Qtran. 2019. Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *Proceedings of the 31st International Conference on Machine Learning, Proceedings of Machine Learning Research*. PMLR.
- [7] Jian Hu, Siyang Jiang, Seth Austin Harding, Haibin Wu, and SW Liao. 2021. Rethinking the implementation tricks and monotonicity constraint in cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2102.03479* (2021).
- [8] Maximilian Hüttenrauch, Adrian Šošić, and Gerhard Neumann. 2017. Guided deep reinforcement learning for swarm systems. *arXiv preprint arXiv:1709.06011* (2017).
- [9] Shariq Iqbal and Fei Sha. 2019. Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2961–2970.
- [10] Kenneth E Kinneer, William B Langdon, Lee Spector, Peter J Angeline, and Una-May O’Reilly. 1994. *Advances in genetic programming*. Vol. 3. MIT press.
- [11] Jinoh Lee, Pyung Hun Chang, and Rodrigo S Jamisola. 2013. Relative impedance control for dual-arm robots performing asymmetric bimanual tasks. *IEEE transactions on industrial electronics* 61, 7 (2013), 3786–3796.
- [12] Jiahui Li, Kun Kuang, Baoxiang Wang, Furui Liu, Long Chen, Fei Wu, and Jun Xiao. 2021. Shapley Counterfactual Credits for Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2106.00285* (2021).
- [13] Shanqi Liu, Licheng Wen, Jinhao Cui, Xuemeng Yang, Junjie Cao, and Yong Liu. [n.d.]. Moving Forward in Formation: A Decentralized Hierarchical Learning Approach to Multi-Agent Moving Together. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 4777–4784.
- [14] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *arXiv preprint arXiv:1706.02275* (2017).
- [15] Anuj Mahajan, Tabish Rashid, Mikayel Samvelyan, and Shimon Whiteson. 2019. Maven: Multi-agent variational exploration. *arXiv preprint arXiv:1910.07483* (2019).
- [16] Junhyuk Oh, Valliappa Chockalingam, Honglak Lee, et al. 2016. Control of memory, active perception, and action in minecraft. In *International Conference on Machine Learning*. PMLR, 2790–2799.
- [17] Liviu Panait, Sean Luke, and R Paul Wiegand. 2006. Biasing coevolutionary search for optimal multiagent behaviors. *IEEE Transactions on Evolutionary Computation* 10, 6 (2006), 629–645.
- [18] Daniel Rakita, Bilge Mutlu, Michael Gleicher, and Laura M Hiatt. 2019. Shared control–based bimanual robot manipulation. *Science Robotics* 4, 30 (2019), eaaw0955.
- [19] Tabish Rashid, Gregory Farquhar, Bei Peng, and Shimon Whiteson. 2020. Weighted qmix: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning. *arXiv preprint arXiv:2006.10800* (2020).
- [20] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2018. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 4295–4304.
- [21] Lloyd S Shapley. 2016. *A value for n-person games*. Princeton University Press.
- [22] 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.
- [23] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Viničius 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).
- [24] Lipeng Wan, Zeyang Liu, Xingyu Chen, Han Wang, and Xuguang Lan. 2021. Greedy-based Value Representation for Optimal Coordination in Multi-agent Reinforcement Learning. *arXiv preprint arXiv:2112.04454* (2021).
- [25] 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).
- [26] Jianhong Wang, Yuan Zhang, Tae-Kyun Kim, and Yunjie Gu. 2020. Shapley Q-value: a local reward approach to solve global reward games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 7285–7292.
- [27] Chenguang Yang, Yiming Jiang, Zhijun Li, Wei He, and Chun-Yi Su. 2016. Neural control of bimanual robots with guaranteed global stability and motion precision. *IEEE Transactions on Industrial Informatics* 13, 3 (2016), 1162–1171.
- [28] Yaodong Yang, Jianye Hao, Guangyong Chen, Hongyao Tang, Yingfeng Chen, Yujing Hu, Changjie Fan, and Zhongyu Wei. 2020. Q-value path decomposition for deep multiagent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 10706–10715.
- [29] Lianmin Zheng, Jiacheng Yang, Han Cai, Ming Zhou, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Magent: A many-agent reinforcement learning platform for artificial collective intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.