ConcaveQ: Non-monotonic Value Function Factorization via Concave Representations in Deep Multi-Agent Reinforcement Learning

Huiqun Li¹, Hanhan Zhou², Yifei Zou^{1*}, Dongxiao Yu¹, Tian Lan^{2*}

¹ Shandong University ² The George Washington University hqleee@outlook.com, {yfzou, dxyu}@sdu.edu.cn, {hanhan, tlan}@gwu.edu

Abstract

Value function factorization has achieved great success in multi-agent reinforcement learning by optimizing joint actionvalue functions through the maximization of factorized peragent utilities. To ensure Individual-Global-Maximum property, existing works often focus on value factorization using monotonic functions, which are known to result in restricted representation expressiveness. In this paper, we analyze the limitations of monotonic factorization and present ConcaveQ, a novel non-monotonic value function factorization approach that goes beyond monotonic mixing functions and employs neural network representations of concave mixing functions. Leveraging the concave property in factorization, an iterative action selection scheme is developed to obtain optimal joint actions during training. It is used to update agents' local policy networks, enabling fully decentralized execution. The effectiveness of the proposed ConcaveQ is validated across scenarios involving multi-agent predator-prey environment and StarCraft II micromanagement tasks. Empirical results exhibit significant improvement of ConcaveO over state-ofthe-art multi-agent reinforcement learning approaches.

Introduction

Joint decision-making in multi-agent reinforcement learning (MARL) often requires dealing with an exponential expansion of the joint state-action space as the number of agents increases. To this end, recent MARL algorithms (Sunehag et al. 2018; Wang et al. 2021a; Rashid et al. 2018; Zhang et al. 2021; Wang et al. 2021b) often leverage value-function factorization techniques – which optimize joint action-value functions through the maximization of factorized, per-agent utilities – to improve learning and sampling efficiency.

Current value factorization methods mostly adhere to the Individual-Global-Max (IGM) (Rashid et al. 2018) principle, which asserts the coherence between joint and individual greedy action selections. It requires value function factorization to be strictly monotonic, i.e., if local agents maximize their local utilities, the joint value function can attain the global maximum. While such monotonic factorization facilitates efficient maximization during training and simplifies the decentralization of the learned policies, it is restricted to monotonic value function representations and may lead to poor performance in situations that exhibit non-monotonic characteristics (Rashid et al. 2020). In response to the limited representation expressiveness due to monotonic factorization, recent works have considered a number of mitigation strategies, e.g., WOMIX (Rashid et al. 2020) that formulates a weighted projection toward monotonic functions, QTRAN (Son et al. 2019) that utilizes an additional advantage value estimator, QPLEX (Wang et al. 2021a) that leverages a dueling structure, and ResQ (Shen et al. 2022) that introduces an additional residual function network. However, these methods tend to compensate for the representation gap caused by monotonic factorization, subject to the IGM principle. QTRAN exhibits low sample efficiency and QPLEX may obtain suboptimal results, while WQMIX tends to possess high approximation errors in non-optimal actions, and it is difficult for ResQ to find the optimal actions for a scenario requiring highly-coordinated agent exploration (Shen et al. 2022). Going beyond monotonicity and IGM while ensuring decentralized action selection is still an open question.

The pivotal insight in this paper is that monotonic valuefunction factorization may not be necessary for enabling local action selections in MARL. Indeed, we show that monotonic factorization under IGM may lead to suboptimal policies that only recover an arbitrarily small fraction of optimal global actions that are achievable through an ideal policy with full observability. Instead, we propose a novel approach of concave value-function factorization for high representation expressiveness, called ConcaveQ. It employs a deep neural network representation of concave mixing networks for value-function decomposition. The concave mixing network estimates the joint action-value output Q_{tot} through a concave function $f_{concave}$ of the input local utilities Q_i conditioned on agent i's local observation and action choice, that is, $Q_{tot} = f_{concave}(Q_1, ..., Q_n)$. We propose a deep neural network representation of the concave mixing function $f_{concave}$. It enables a novel concave value-function factorization in MARL.

Due to the lack of IGM under concave mixing functions, the global optimal joint actions are not necessarily consistent with the optimal actions of local agents. ConcaveQ tackles this problem through an iterative action-selection algorithm during training, which attains optimality using the concavity property of the proposed value-function factorization. To sup-

^{*}Yifei Zou and Tian Lan are corresponding authors.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

port fully decentralized execution, we note that the optimal joint action cannot be obtained directly from maximizing the local agent utilities Q_i . ConcaveQ adopts a soft actor-critic structure with local policy networks for distributed execution, such that each local agent can select the best action according to its own local policy network, which are corrected by the concave value networks during training. It also allows auxiliary information such as entropy maximization to be exploited for effective learning.

For evaluation, we demonstrate the performance of ConcaveQ on SMAC maps and predator-prey tasks. We compare ConcaveQ with several state-of-the-art value factorization methods on various MARL datasets. The experimental results show that ConcaveQ outperforms multiple competitive value factorization methods, especially on difficult tasks that require more coordination among the agents since monotonic factorization restricts global value-function estimate and hinders the effective training of agents. Our numerical results show significant improvement over SOTA baselines in Star-Craft II micromanagement challenge tasks regarding higher reward and convergence speed.

Background

Value Decomposition. Value function decomposition is a paradigm for cooperative multi-agent reinforcement learning (MARL) that aims to learn a centralized state-action value function that can be decomposed into individual agent values. Value decomposition approaches can reduce the complexity and improve the scalability of MARL problems. Two prominent examples of value function decomposition methods are VDN and QMIX, which both assume that the centralized action value Q_{tot} is additive or monotonic with respect to the agent values Q_i . VDN simply sums up the agent values to obtain $Q_{tot} = \sum_{i=1}^{n} Q_i(\tau_i, u_i)$, while QMIX uses a mixing network to combine the local agent utilities in a monotonic way. QMIX enforces the monotonicity through a constraint on the relationship between $Q_{tot}(s, \boldsymbol{u}) = f_{\theta}(s, Q_1(\tau_1, u_1), ..., Q_n(\tau_n, u_n))$ and each Q_i , that is, $\frac{\partial Q_{tot}(\boldsymbol{\tau}, \boldsymbol{u})}{\partial Q_i(\tau_i, u_i)} > 0, \forall i \in \mathcal{N}$. The monotonicity constraint ensures that the optimal joint action for the global reward Q_{tot} is also optimal for each agent's local utility Q_i , where the mixing function f_{θ} is formulated as a feed-forward network parameterized by θ . The weights of the mixing network are produced by independent hyper-networks, which take the global state as input and use an absolute activation function to ensure that the mixing network weights are non-negative to enforce the monotonicity. Then QMIX is trained end-to-end to minimize the squared TD error on mini-batch of b samples from the replay buffer as $\Sigma_{i=1}^{b} (Q_{tot}(\boldsymbol{\tau}, \boldsymbol{u}, s; \theta) - y_{tot})^2$, where $y^{tot} = r + \gamma max_{u'}Q_{tot}(\boldsymbol{\tau'}, \boldsymbol{u'}, s'; \theta')$, r is the global reward and θ' is the parameters of the target network whose parameters are periodically copied from θ for training stabilization.

Weighted QMIX Projection. QMIX (Rashid et al. 2018) projects Q_{tot} to the restricted function class Q^{mix} by mini-

mizing the projection loss:

$$\underset{q \in \mathcal{Q}^{mix}}{\operatorname{argmin}} \sum_{\mathbf{u} \in \mathbf{U}} (\mathcal{T}^* Q_{tot}(s, \mathbf{u}) - q(s, \mathbf{u}))^2.$$
(1)

Weighted QMIX (Rashid et al. 2020) extends the idea of monotonic mixing by introducing a weighting function into the QMIX projection operator to bias the learning process towards the best joint action. The weighted projection can be described as:

$$\underset{q \in Q^{mix}}{\operatorname{arg\,mix}} \sum_{\mathbf{u} \in \mathbf{U}} (w(s, \mathbf{u})Q(s, \mathbf{u}) - q(s, \mathbf{u}))^2$$
(2)

where w is the weighting function that is added to place more importance on optimal joint actions, while still anchoring down the value estimates for other joint actions. Weighted QMIX can be viewed as a projection onto monotonic mixing function space using weighted distance.

Related Work

Value Factorization Approaches. Value Factorization approaches are widely adopted in value-based MARL (Hu et al. 2021; Wang et al. 2021b; Zhou, Lan, and Aggarwal 2023; Gogineni et al. 2023; Mei et al. 2023). Current value factorization methods mostly adhere to the monotonic IGM principle, such as VDN(Sunehag et al. 2018) and QMIX(Rashid et al. 2018). VDN (Sunehag et al. 2018) represents the joint state-action value function Q_{tot} as a sum of per-agent utilities Q_i . QMIX(Rashid et al. 2018) employs a mixing network to factorize the Q_{tot} in a monotonic manner. The monotonicity constraint ensures that the optimal joint action for the global reward Q_{tot} is also optimal for each agent's local utility Q_i . However, these two decomposition methods suffer from structural constraints, limiting the range of joint action-value functions they can effectively represent.

To compensate for the restricted expressiveness of monotonic representation, recent works have explored several mitigation strategies. Specifically, WQMIX (Rashid et al. 2020) applies a weighted projection to QMIX, which attaches more importance to the optimal joint actions when minimizing training errors. In WQMIX, finding the optimal weight remains an open problem. Furthermore, the approximation errors are high for non-optimal actions. QTRAN (Son et al. 2019) incorporates an additional advantage value estimator and imposes a set of linear constraints. QPLEX (Wang et al. 2021a) leverages a dueling structure involving value and advantage functions. ResQ (Shen et al. 2022) masks some state-action value pairs and introduces an additional residual function network. Apart from mitigation strategies, there are also various MARL actor-critic methods, such as MAD-DPG (Lowe et al. 2017), MAAC (Iqbal and Sha 2019), and COMA (Foerster et al. 2018), which use centralized critics and decentralized actors. PAC (Zhou, Lan, and Aggarwal 2022) decouples individual agents' policy networks from value function networks to leverage the benefits of assisted value function factorization. VDAC (Su, Adams, and Beling 2021) combined actor-critic structure with QMIX for the joint state-value function estimation. DOP (Wang et al. 2021b) employs a network similar to Qatten (Yang et al. 2020) for policy gradients with off-policy tree backup and on-policy TD.

From the previous works mentioned above, we can see a series of works conducted within the confines of the monotonic assumption. However, few of them consider the nonmonotonic case. To the best of our knowledge, this paper is the first one considering the MARL in the concave scenario, which is a general case in non-monotonic scenarios.

Characterizing Limitations of Monotonic Factorization

Monotonic factorization restricts the family of global action value functions that can be represented. More precisely, $Q_{tot}(\mathbf{s}, \boldsymbol{u}) = f_{\theta}(\mathbf{s}, Q_1(\tau_1, u_1), ..., Q_n(\tau_n, u_n))$ implies that if $u_1^* = \operatorname{argmax} Q_1(\tau_1, u_1)$ is the optimal action choice for agent 1, u_1^* must also be optimal for all other states with varying τ_2, \ldots, τ_n . We analyze this restriction and show how it makes monotonic factorization ineffective in representing complex global value functions.

Considering a Markov decision process (MDP) with nagents, where the state of agent i is denoted by $s_i \in S$ and action by $u_i \in A$, for i = 1, 2, ..., n, composing a joint state **s** and joint action **u**. Let $Q_i(s_i, u_i)$ be local agent value utilities, $Q_{jt}(\mathbf{s}, \mathbf{u})$ be the unrestricted ground truth of joint actionvalue function, and $Q_{mono}(\mathbf{s}, \mathbf{u})$ be an arbitrary monotonic factorization estimate of $Q_{jt}(\mathbf{s}, \mathbf{u})$. Ideally, the monotonic factorization estimate Q_{mono} should be able to recover the exact optimal action selections of Q_{jt} . We define S_{mono} as the state set, where Q_{mono} and Q_{jt} have the same optimal joint action, i.e.,

$$S_{mono} = \{\mathbf{s} : \arg\max Q_{mono}(\mathbf{s}, \mathbf{u}) = \arg\max Q_{jt}(\mathbf{s}, \mathbf{u})\}$$

We show that $|S_{mono}|$ can be arbitrarily small compared to the state space $|S|^n$, when the global maximums are uniformly distributed in $Q_{jt}(\mathbf{s}, \mathbf{u})$. Thus, monotonic factorization could only recover an arbitrarily small fraction of action choices.

Theorem 1 When $n \geq \log_{|S|}(2|A| \cdot \log_2|A|) + 1$ and the optimal action choices in Q_{jt} are uniformly distributed, for any monotonic factorization, we have $\frac{E(|S_{mono}|)}{|S|^n} \leq \frac{e+1}{|A|} \cdot \delta^{n-1}$ for some constant $\delta \in (0, 1)$.

Proof Sketch. We give a sketch of the proof and provide the complete proof in Appendices. Our key idea is to convert the problem into a classic max-load problem (Czumaj 2004).

Step1: Formulate as max-load bin-ball problem. For each agent *i* and state s, we consider the optimal action of Q_{mono} as ball *i*. Thus, Q_{jt} and Q_{mono} have the same optimal action for agent *i* if ball *i* is placed in the bin corresponding to the optimal action of Q_{jt} .

Let X_i denotes that ball *i* is in bin *i*, that is, Q_{jt} and Q_{mono} have the same optimal action for agent *i*, then we have:

$$E(|S_{mono}|) = \Sigma_s P(X_1) \cdot P(X_2|X_1) \cdot \dots \cdot P(X_n|X_{n-1} \dots X_1)$$

Define Y_i as the load of bin *i*, that is, the state space such that Q_{jt} and Q_{mono} have the same optimal action for agents except for agent *i*. Note that the global maximum of Q_{jt} is uniformly distributed over different states, we then analyze $P(X_1)$:

$$P(X_1) = \frac{E(\sum_{s_2...s_n} X_1)}{|S|^{n-1}} \le \frac{E(max_i Y_i)}{|S|^{n-1}}$$
(3)

Step2: Analyze the probability distribution of the load. According to the Chernoff bound and the Union bound, we have: when $n \ge log_{|S|}(2|A| \cdot log_2|A|) + 1$,

$$E(max_iY_i) \le \frac{(e+1) \cdot |S|^{n-1}}{|A|}.$$
 (4)

where $|A| \ge 1$ is the size of action space.

Applying Eq. (4) to Eq. (3) and Eq. (3), we have:

$$P(X_1) \le \frac{e+1}{|A|} \tag{5}$$

Using the same argument repeatedly for agents i = 2, ..., n, we can choose $\delta = min_i(P(X_i|X_{i-1}...X_1))$ and $0 < \delta < 1$. Plugging these inequalities into $E(|S_{mono}|)$, it yields the desired result $\frac{E(|S_{mono}|)}{|S|^n} \leq \frac{e+1}{|A|} \cdot \delta^{n-1}$. Our analysis shows that a monotonic mixing can only re-

cover an arbitrarily small fraction of global value function maximums, as the number of agents or the size of action space increases. Our proposed ConcaveQ addresses the representation issue by introducing a concave mixing network to value function factorization, to enhance the representation ability for better performance. Since the IGM principle no longer holds under concave mixing functions, the greedy action selection by maximizing local agent Q values can not be adopted. Note that concave optimization has good convergence properties (Karush 1939), that is, a concave function has only one global maximum and the maximum can always be obtained using iterative methods. Therefore, ConcaveQ tackles the action selection problem by optimizing the joint value function in an iterative manner. As for execution, note that the optimal joint action cannot be obtained directly from maximizing the local agent Q_i , we adopt a soft actor-critic policy network in ConcaveQ that uses factorized policies to support distributed execution, such that each local agent can select the best action according to its own local policy network which will be corrected by the Concave value networks during training while exploiting auxiliary information for learning. Moreover, entropy maximization is also used to enhance effective exploration. As for concave mixing networks, there is always a concave mixing way that can recover the global maximum.

Proposition 1 For any state s, there is always a concave function $Q_c(s, \mathbf{u})$ that recovers the global maximum of $Q_{jt}(s, \mathbf{u})$ with the same optimal action.

Proof Sketch. For any state s, let $f_1(\mathbf{u}) = Q_{jt}(\mathbf{u})$ and $f_2(\mathbf{u}) = -Q_{jt}(\mathbf{u})$. According to Fenchel–Moreau theorem (Borwein and Lewis 2006), the biconjugate function of $f_2(\mathbf{u})$ is a convex function $g(\mathbf{u}) = f_2^{**}(\mathbf{u})$ and $h(\mathbf{u}) = -f_2^{**}(\mathbf{u})$ is a concave function. $g(\mathbf{u})$ and $h(\mathbf{u})$ satisfies:

$$g(\mathbf{u}) \le f_2(\mathbf{u}),\tag{6}$$

$$h(\mathbf{u}) \ge f_1(\mathbf{u}). \tag{7}$$

Suppose the optimal joint action of $f_1(\mathbf{u})$ and $h(\mathbf{u})$ are $\mathbf{u}_{f_1}^*$ and \mathbf{u}_h^* respectively, if we shift $h(\mathbf{u})$ by $(-\mathbf{u}_{f_1}^* + \mathbf{u}_h^*)$, the



Figure 1: The overall architecture of ConcaveQ. The concave mixing network represents Q_{tot} as a concave function of local Q_a . With the help of iterative action selection, we can select the optimal actions during training. As for the execution process, the soft policy network is used to find the optimal actions.

shifted concave function has the same optimal action as $f_1(\mathbf{u})$. In other words, for any state s, there is always a concave function that has the same optimal action as $Q_{jt}(s, \mathbf{u})$.

The key to our method is the insight that it is not necessary to use the monotonic factorization of QMIX to extract decentralized policies that are fully consistent with their centralized counterpart. Next, we will discuss concave mixing network design and our ConcaveQ framework to support fully decentralized execution.

Method

Since concave factorization may not satisfy the IGM principle, we introduce an iterative action selection algorithm to ensure optimal joint action selection during training. It iteratively optimizes each agent's action selection with all other agent's actions and local utilities fixed. Due to the concave property of the mixing function, the iterative algorithm converges to global optimal actions. Note that this iterative action selection does not directly allow decentralized execution. To this end, we further leverage the soft actor-critic framework to train local factorized policies for distributed execution. Figure 1 shows the architecture of our learning framework. There are four main components in ConcaveQ : (1) a concave Q_{tot} utilizing per-agent local utilities, where a concave function serves as the mixing network (2) an unrestricted joint action estimator \hat{Q}^* , which serves as a baseline estimator of the true optimal value function Q^* (3) an iterative action-selection module to seek for the optimal joint action during training, and (4) a soft-policy network that allows for the utilization of a soft actor-critic network that consists of local policy networks and local value networks to perform decentralized execution.

Concave Mixing Network Design

The concave mixing network constrains the output of the network to be a concave function of the inputs, which covers a large class of functions and allows efficient inference via optimization over inputs to the network. Moreover, the concave feature ensures that maximizing the output of the network leads to only one solution. This means that theoretically there won't be local maximum or multiple solutions that may confuse the inference process. Concave mixing network has a suitable property specifically for value function factorizationbased multi-agent reinforcement learning problems, that is, local maximization equals global maximization and it can be obtained through iterative methods, which is of great use in our case.

We first consider a k-layer, fully connected network as shown in the left part of Figure 1. A concave neural network is defined over the input x for i = 0, ..., k - 1 using such architecture:

$$z_{i+1} = \begin{cases} W_0^{(z)} x + b_0, i = 0\\ r_i(W_i^{(z)}[z_1, ..., z_i] + b_i), i = 1, ..., k - 2\\ -W_i^{(z)}[z_1, ..., z_i] + b_i, i = k - 1 \end{cases}$$
(8)

where z_i denotes the layer activations, r_i are non-linear activation functions, and $\theta = (W_{0:k-1}^{(z)}, b_{0:k-1})$ are the parameters of linear layers. We can model the concave mixing network as:

$$f(x;\theta) = z_k \tag{9}$$

Theorem 2 The mixing network f is concave in x provided that all $W_{1:k-1}^{(z)}$ are non-negative, and all functions r_i are convex and non-decreasing.

We delineate the proof and the detailed proof is provided in Appendices. The proof is simple and follows from the fact that non-negative sums of convex functions are also convex and that the composition of a convex and convex non-decreasing function is also convex, and that the negative of a convex function is a concave one (Boyd and Vandenberghe 2004). In our framework, we satisfy the constraint that the r_i should be convex non-decreasing via nonlinear activation units, specifically we adopt the rectified linear functions. The constraint that the $W^{(z)}$ terms be non-negative is somewhat restrictive, but because the bias terms can be negative, the network still has substantial representation power. In our multi-agent reinforcement learning setting, we model $Q_{tot}(s, \mathbf{u}; \theta)$ using an input-concave mixing function and select actions under the concave optimization problem $u^*(s) = \operatorname{argmax}_u Q_{tot}(s, \mathbf{u}; \theta)$. We have 4 layers, k = 4 and the weights $W_{0:k-1}^{(z)}$ are generated by hypernetworks. Each hypernetwork takes the state s as input and consists of a single linear layer, succeeded by an absolute ReLU activation function, to generate the non-negative weights.

Iterative Action-Selection and Local Policies

We note that the optimal joint action of Q_{tot} cannot be obtained directly from maximizing Q_i since the mixing function is concave rather than monotonic. Hence it is almost impossible to directly find the maximum Q_{tot} and the corresponding joint action selections \mathbf{u}^* such that $\mathbf{u}^* = \operatorname{argmax}_{\mathbf{u}}\mathbb{E}[Q(s, \mathbf{u})]$, as this takes $O(|U|^{|n|})$ iterations to find the maximum. One of the key insights underlying this method is that although it is impractical to find \mathbf{u}^* directly: however, its local estimation and local optimum $\mathbf{u}^* = \operatorname{argmax}_{\mathbf{u}}\mathbb{E}[Q^*(s, \tau, \cdot)]$, which takes O(|U| * |n|) time, is possible to find which will update the concave mixing network and local policy network asymptotically.

To illustrate this problem, we consider the comparison between the action selection from the monotonic mixing network and concave mixing network: in value-based methods with a monotonic mixing network, each agent chooses $u_t^i = \operatorname{argmax}(q^i(\tau_t^i))$, i.e. each agent chooses the action with the best local value thus getting an optimal global value; however, when the utilizing a concave mixing network for a more general value function factorization, we can maximize $\mathbb{E}[Q^*(s, \tau, \cdot)]$ via iterative action selection.

Next, since the mixing network is non-monotonic (concave), we can no longer follow the IGM principle to adopt actions that maximize individual values from the local value networks. Instead, we adopt a soft actor-critic policy network that uses factorized policies to enable distributed execution, such that each local agent can choose the best action according to its local policy network which will be corrected by the concave value networks during training while exploiting auxiliary information for learning. Under a soft-actor-critic paradigm, each agent can choose $u_t^i = \operatorname{argmax}(\pi^i(\tau_t^i))$ such that $\mathbb{E}[Q^*(s,\tau,\mathbf{u})] = \max(\mathbb{E}[Q^*(s,\tau,\cdot)])$. This ensures that the concave mixing network could factorize concave value functions while each local agent could choose their best actions based on local policy in a decentralized manner during execution. Previous studies have demonstrated that Boltzmann exploration policy iteration can guarantee policy improvement and optimal convergence with infinite iterations and complete policy evaluation. With factorized policy, we have local policy trained with the following:

$$\mathcal{L}_{\pi}(\theta) = \mathbb{E}_{\mathcal{D}} \left[\boldsymbol{\alpha} \log \boldsymbol{\pi} \left(\boldsymbol{u}_{t} | \boldsymbol{\tau}_{t} \right) - Q_{tot}^{\pi} \left(\boldsymbol{s}_{t}, \boldsymbol{\tau}_{t}, \boldsymbol{u}_{t} \right) \right]$$
$$= -q^{\pi} \left(\boldsymbol{s}_{t}, \mathbb{E}_{\pi^{i}} \left[q^{i} \left(\boldsymbol{\tau}_{t}^{i}, \boldsymbol{u}_{t}^{i} \right) - \alpha^{i} \log \pi^{i} \left(\boldsymbol{u}_{t}^{i} | \boldsymbol{\tau}_{t}^{i} \right) \right] \right)$$
(10)

where q^{π} is the local value network with $u_i \sim \pi_i(o_i)$, \mathcal{D} is the replay buffer for sampling training data, and the parameter α from soft-actor-critic controls how much the agents prefer actions with higher expected rewards over actions with more explorations. We introduce the training of the local value network and other components in the following sections. The algorithm has a more vital exploration ability and a higher level of generalization.

Training ConcaveQ

We have presented the components of our approach so far. Next, we describe how to implement and train our novel RL algorithm that scales well under DEC-POMDP. During the decentralized execution phase, only the agent network (local policy in Fig.1) is active with actions chosen greedily from each local policy network as $u_t^i = \operatorname{argmax}(\pi^i(\tau_t^i))$, ensuring full CTDE.

The \hat{Q}^* architecture (Green part in Fig. 1) is used as the estimator for Q^* from unrestricted functions, where it serves as a mixing network using a feed-forward network that takes its local utilities. Together with the proposed concave mixing network, they are trained to reduce the following loss:

$$\mathcal{L}_{\hat{Q^*}}(\theta) = \sum_i (\hat{Q^*}(s, \boldsymbol{\tau}, \hat{\boldsymbol{u}}) - y_i)^2 \tag{11}$$

$$\mathcal{L}_{\text{ConcaveQ}}(\theta) = \sum_{i} w(s, \mathbf{u}) (Q_{tot}(s, \tau, \mathbf{u}) - y_i)^2 \quad (12)$$

where $\hat{\boldsymbol{u}} = \operatorname{argmax}_{\hat{\boldsymbol{u}}} Q_{tot}(\boldsymbol{\tau}', \hat{\boldsymbol{u}}', s'; \boldsymbol{\theta})$ is from local iterative action selection, $y_i = r + \gamma \hat{Q}^*(s', \boldsymbol{\tau}', \hat{\boldsymbol{u}})$ and $\boldsymbol{\theta}'$ is the parameters of the target network that are periodically updated to stabilize the training. $w(s, \boldsymbol{u})$ is the weighting function with w = 1 if $Q_{tot}(s, \tau, \boldsymbol{u}) - y_i < 0, w = 0.5$ otherwise as suggested in WQMIX (Rashid et al. 2020). Then all components are trained in an end-to-end manner as:

$$\mathcal{L}(\theta) = \mathcal{L}_{\pi} + \mathcal{L}_{\hat{O}^*} + \mathcal{L}_{\text{ConcaveQ}}$$
(13)

We provide detailed derivations and pseudo-code for the training process in the appendices.

Experiment Results

In this section, we present our experimental results with the state-of-the-art methods on Predator-Prey and the StarCraft II Multi-Agent Challenge (SMAC)(Samvelyan et al. 2019). For fair evaluations, the hyper-parameters of all algorithms under comparison as well as the optimizers are the same. Each experiment is repeated 3 times with different seeds. The presented curves are smoothed by a moving average filter with its window size set to 5 for better visualization. More implementation details and experimental introduction and settings can be found in the Appendices.



Figure 2: Average reward on the Predator-Prey tasks.



Figure 3: Average test win rate on the SMAC tasks.

Predator Prey

Predator-Prey is a complex partially observable multi-agent environment, where 8 agents cooperate as predators to hunt 8 prey within a 10×10 grid. If two or more adjacent predator agents carry out the *catch* action simultaneously, it is a successful catch and the agents receive a reward r = 10. A failed attempt where only one agent captures the prey will receive a punishment reward $p \leq 0$. The more negative p is, the higher level of coordination is needed for the agents.

We select multiple state-of-the-art MARL approaches as baseline algorithms for comparison, which includes valuebased factorization algorithms (i.e., QMIX in (Rashid et al. 2018), WQMIX in (Rashid et al. 2020), PAC in (Zhou, Lan, and Aggarwal 2022), QPLEX in (Wang et al. 2021a), and RESO in (Shen et al. 2022)), and decomposed actor-critic approaches (i.e., FOP (Zhang et al. 2021)). Fig 2 shows the performance of our scheme and the six baselines when punishment p varies from 0 to -2, in which the x-axes and y-axes represent the number of training episodes and the test mean rewards, respectively. According to the curves in Fig 2, the following results can be observed. (1) When p = 0, CON-CAVEQ, PAC, WQMIX, and QMIX can learn good policies and obtain the highest reward when they have been trained for more than 50 episodes. In other words, the performance of our algorithm is as good as the state-of-the-art works. (2) When punishment gets larger, i.e. p = -0.5, p = -1.5 and p = -2.0, only CONCAVEQ and WQMIX can still achieve high rewards within 60 episodes while CONCAVEO is able to converge faster than WQMIX, the others gradually fail due to their monotonicity constraints or representational limits. These results demonstrate CONCAVEQ's ability in challenging cooperative MARL tasks that require non-monotonic action selections.

StarCraft II Multi-Agent Challenge (SMAC)

In SMAC, agents are divided into two teams to cooperate with allies and compete against enemies or against the other team controlled by the built-in game AI. In the simulation, the agents act according to their local observations and learning experiences.

Note that the state-of-the-art algorithms have already achieved a very good performance on the easy and medium maps, which makes it difficult to present clear comparisons and potential improvements. We carry out our experiment in six maps consisting of two hard maps (3s_vs_5z, 5m_vs_6m) and four super-hard maps (27m_vs_30m, 6h_vs_8z, MMM2, corridor). The baseline algorithms are the same as those in the Predator-Prey environment.

Details of the environment setting and other training details like network hyperparameters can be found in Appendices.

The performance of our algorithm and baselines in hard and super hard maps are presented in Figure 3, in which the *x*-axes and *y*-axes represent the number of training episodes and the test win rate, respectively. For almost all scenarios, we found that our algorithm is able to converge faster and deliver a higher win rate. Especially on the $6h_{VS}_{8z}$ map, CONCAVEQ significantly outperforms other baselines by a large margin. We note that the $6h_{VS}_{8z}$ map requires a sophisticated strategy to take the win, where one unit draws firepower from the enemy units while circling around the edge of the map so the rest friendly units can deliver damages. This implies CONCAVEQ imposes a higher capability of exploration and function representational abilities.



Figure 4: Ablations results comparing CONCAVEQ and its ablated versions on SMAC map 3s_vs_5z

Ablation Studies

We carry out ablation experiments to demonstrate the effectiveness and contribution of each core component introduced in CONCAVEQ on 3s_vs_5z scenario in SMAC. As shown in Fig.4 we consider verifying the effect of (1) concave mixing network by replacing it with 4-layer monotonic network as CONCAVEQ_linear, (2) iterative action selection which removes the iterative procedure as CONCAVEQ_no_iter, (3) soft policy network which removes the soft policy network as CONCAVEQ_no_policy, (4) disable both iterative action selection and soft policy network as CONCAVEQ_disabled, (5) further remove the central Q^* as $CONCAVE_no_Q^*$. The results show that CONCAVEQ overperforms these ablated versions. CONCAVEQ_no_iter has lower rewards than CON-CAVEQ due to a lack of iterative action selection which is unfavorable to finding the optimal actions. Removing the soft policy network also leads to performance drop due to lack of proper action selection scheme during execution. The reward of CONCAVEO_linear grows slowly at first as linear layers have weaker representation ability than concave mixing networks. $CONCAVEQ_no_Q^*$ suffers from the most significant performance drop after most core components are removed from the original design. Such results validate how each component is crucial for achieving performance through experiments.

Conclusions

In this paper, we propose ConcaveQ, a novel non-monotonic value function factorization approach, which goes beyond monotonic mixing functions that are known to have limited representation expressiveness. ConcaveQ employs neural network representations of concave mixing functions. An iterative action selection scheme is developed to obtain optimal action during training, while factorized policies using local networks are leveraged to support fully decentralized execution. We evaluate the proposed ConcaveQ on predator-prey and StarCraft II tasks. Empirical results demonstrate substantial improvement of ConcaveQ over state-of-the-art MARL algorithms with monotonic factorization.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (NSFC) under Grant 62122042 and in part by Major Basic Research Program of Shandong Provincial Natural Science Foundation under Grant ZR2022ZD02.

References

Borwein, J.; and Lewis, A. 2006. *Convex Analysis and Nonlinear Optimization: Theory and Examples.* Springer.

Boyd, S.; and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge university press.

Czumaj, A. 2004. Lecture notes in Approximation and Randomized Algorithms. https://www.ic.unicamp.br/~celio/ peer2peer/math/czumaj-balls-into-bins.pdf. Accessed: 2023-02-01.

Foerster, J. N.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual Multi-Agent Policy Gradients. In *AAAI*.

Gogineni, K.; Mei, Y.; Lan, T.; Wei, P.; and Venkataramani, G. 2023. Accmer: Accelerating multi-agent experience replay with cache locality-aware prioritization. In 2023 IEEE 34th International Conference on Application-specific Systems, Architectures and Processors (ASAP), 205–212. IEEE.

Hu, J.; Jiang, S.; Harding, S. A.; Wu, H.; and Liao, S.-w. 2021. Rethinking the implementation tricks and monotonicity constraint in cooperative multi-agent reinforcement learning. In *ICLR*.

Iqbal, S.; and Sha, F. 2019. Actor-Attention-Critic for Multi-Agent Reinforcement Learning. In *ICML*.

Karush, W. 1939. *Minima of Functions of Several Variables with Inequalities as Side Conditions*. Master's thesis, Department of Mathematics, University of Chicago.

Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *NeuralIPS*.

Mei, Y.; Zhou, H.; Lan, T.; Venkataramani, G.; and Wei, P. 2023. MAC-PO: Multi-Agent Experience Replay via Collective Priority Optimization. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, 466–475. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321.

Rashid, T.; Farquhar, G.; Peng, B.; and Whiteson, S. 2020. Weighted QMIX: Expanding Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *NeuralIPS*.

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 *ICML*.

Samvelyan, M.; Rashid, T.; De Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*. Shen, S.; Qiu, M.; Liu, J.; Liu, W.; Fu, Y.; Liu, X.; and Wang, C. 2022. ResQ: A Residual Q Function-based Approach for Multi-Agent Reinforcement Learning Value Factorization. In *NeurIPS*.

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 *ICML*.

Su, J.; Adams, S.; and Beling, P. 2021. Value-decomposition multi-agent actor-critics. In *AAAI*.

Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W. M.; Zambaldi, V. F.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; and Graepel, T. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *AAMAS*.

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

Wang, Y.; Han, B.; Wang, T.; Dong, H.; and Zhang, C. 2021b. DOP: Off-Policy Multi-Agent Decomposed Policy Gradients. In *ICLR*.

Yang, Y.; Hao, J.; Liao, B.; Shao, K.; Chen, G.; Liu, W.; and Tang, H. 2020. Qatten: A general framework for cooperative multiagent reinforcement learning. *arXiv preprint arXiv:2002.03939*.

Zhang, T.; Li, Y.; Wang, C.; Xie, G.; and Lu, Z. 2021. FOP: Factorizing Optimal Joint Policy of Maximum-Entropy Multi-Agent Reinforcement Learning. In *ICML*.

Zhou, H.; Lan, T.; and Aggarwal, V. 2022. PAC: Assisted Value Factorization with Counterfactual Predictions in Multi-Agent Reinforcement Learning. In *NeurIIPS*.

Zhou, H.; Lan, T.; and Aggarwal, V. 2023. Value Functions Factorization With Latent State Information Sharing in Decentralized Multi-Agent Policy Gradients. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 7(5): 1351–1361.