

A Unified Framework for Factorizing Distributional Value Functions for Multi-Agent Reinforcement Learning

Wei-Fang Sun*

Elsa Lab, Department of Computer Science

National Tsing Hua University

No. 101, Section 2, Kuang-Fu Road, Hsinchu City 30013, Taiwan

J3SOON@GAPP.NTHU.EDU.TW

Cheng-Kuang Lee

NVIDIA AI Technology Center

NVIDIA Corporation

Santa Clara, CA 94305 USA

CKL@NVIDIA.COM

Simon See

NVIDIA AI Technology Center

NVIDIA Corporation

Santa Clara, CA 94305 USA

SSEE@NVIDIA.COM

Chun-Yi Lee†

Elsa Lab, Department of Computer Science

National Tsing Hua University

No. 101, Section 2, Kuang-Fu Road, Hsinchu City 30013, Taiwan

CYLEE@CS.NTHU.EDU.TW

Editor: Shipra Agrawal

Abstract

In fully cooperative multi-agent reinforcement learning (MARL) settings, environments are highly stochastic due to the partial observability of each agent and the continuously changing policies of other agents. To address the above issues, we proposed a unified framework, called DFAC, for integrating distributional RL with value function factorization methods. This framework generalizes expected value function factorization methods to enable the factorization of return distributions. To validate DFAC, we first demonstrate its ability to factorize the value functions of a simple matrix game with stochastic rewards. Then, we perform experiments on all *Super Hard* maps of the StarCraft Multi-Agent Challenge and six self-designed *Ultra Hard* maps, showing that DFAC is able to outperform a number of baselines.

Keywords: Reinforcement Learning, Multi-Agent RL, Distributional RL, Value Function Factorization

1. Introduction

In multi-agent reinforcement learning (MARL), one of the popular research directions is to enhance the training procedure of fully cooperative and decentralized agents. Examples of such agents include a fleet of unmanned aerial vehicles (UAVs), a group of autonomous cars,

*. Wei-Fang Sun contributed to the work during his NVIDIA internship.

†. Corresponding Author.

etc. This research direction aims to develop a decentralized and cooperative behavior policy for each agent, and is especially difficult for MARL settings without an explicit communication channel. The most straightforward approach is independent Q-learning (IQL) (Tan, 1993), where each agent is trained independently with its own behavior policy aimed to optimize the total return. Nevertheless, each agent’s policy may not converge owing to two main difficulties: (1) non-stationary environments caused by the changing behaviors of other agents, and (2) spurious reward signals originated from the actions of other agents. In addition, an agent’s partial observability of an environment further exacerbates the above issues. On the other hand, another line of research focuses on fully utilizing all agents’ information to learn a joint cooperative policy with a single monolithic network. Such methods, named fully-centralized learners, eliminate the policy convergence issue. Unfortunately, they are not scalable and are unable to be executed in a decentralized fashion. In order to strike a balance between these two types of approaches, a number of MARL researchers turned their attention to centralized training with decentralized execution (CTDE) methods, with an objective to stabilize the training procedure while maintaining the agents’ abilities for decentralized execution (Oliehoek and Amato, 2016). Among these CTDE approaches, value function factorization methods (Sunehag et al., 2018; Rashid et al., 2018; Son et al., 2019) are especially promising in terms of their superior performances and data efficiency (Samvelyan et al., 2019).

Value function factorization methods aim to learn a factorization of a certain joint value function based on all agents’ information during centralized training, so as to enable decentralized execution based on the learned factorization. These methods introduce the assumption of individual-global-max (IGM) (Son et al., 2019), which assumes that each agent’s optimal actions result in the optimal joint actions of the entire group. Based on IGM, the total return of a group of agents can be factorized into separate utility functions (Guestrin et al., 2001) (or simply ‘*utility*’ hereafter) for each agent. The utilities allow the agents to independently derive their own optimal actions during execution, and deliver promising performance in StarCraft Multi-Agent Challenge (SMAC) (Samvelyan et al., 2019). Unfortunately, current value function factorization methods only concentrate on estimating the expectations of the utilities, overlooking the additional information contained in the full return distributions. Such information, nevertheless, has been demonstrated beneficial for policy learning in the recent literature (Lyle et al., 2019).

In the past few years, distributional RL has been empirically shown to enhance value function estimation in various single-agent RL (SARL) domains (Bellemare et al., 2017; Dabney et al., 2018b,a; Rowland et al., 2019; Yang et al., 2019). Instead of estimating a single scalar Q-value, it approximates the probability distribution of a total return by either a categorical distribution (Bellemare et al., 2017) or a quantile function (Dabney et al., 2018b,a). Even though the above methods may be beneficial to the MARL domain due to the ability to capture uncertainty, it is inherently incompatible to expected value function factorization methods (e.g., value decomposition network (VDN) (Sunehag et al., 2018), monotonic mixing network (QMIX) (Rashid et al., 2018), and duplex dueling mixing network (QPLEX) (Wang et al., 2021b)). The incompatibility arises from two aspects: (1) maintaining IGM in a distributional form, and (2) factorizing the probability distribution of a total return into individual utilities. As a result, an effective and efficient approach that

is able to solve the incompatibility is crucial and necessary for bridging the gap between value function factorization methods and distributional RL.

In this paper, we propose a **D**istributional Value Function **F**actorization (DFAC) framework, to efficiently integrate value function factorization methods with distributional RL. DFAC solves the incompatibility by two techniques: (1) Mean-Shape Decomposition and (2) Quantile Mixture. The former allows the generalization of expected value function factorization methods (e.g., VDN, QMIX, and QPLEX) to their DFAC variants without violating IGM. The latter allows a total return distribution to be factorized into individual utility distributions in a computationally efficient manner. We slightly abuse the term ‘factorizing return distributions’ here, which refer to factorizing the value function into individual utilities, and should not be confused with factorizing a joint distribution into a product of marginal distributions. To validate the effectiveness of DFAC, we first demonstrate the ability of factorizing the return distribution of a matrix game with stochastic rewards. Then, we perform experiments on all *Super Hard* maps in SMAC, and show that DFAC offers beneficial impacts on the baseline methods. Furthermore, we designed six *Ultra Hard* maps to further validate the effectiveness of DFAC. In summary, the primary contribution is the introduction of DFAC for bridging the gap between distributional RL and value function factorization methods efficiently by mean-shape decomposition and quantile mixture. Please note that this work is an extended version of Sun et al. (2021).

2. Related Works

The related works of this paper can be grouped into four categories as follows.

Distributional Independent Learners. Independent learners allow agents to learn a cooperative policy based on their own action-observation histories, without assuming additional information from other agents. The major drawback of this approach is the miscoordination among agents due to the lack of information exchange during training. Several works incorporate distributional RL methods to mitigate this issue. Da Silva et al. (2019) extended the simplest IQL (Tan, 1993), which simply performs expected Q-learning for each agent, to multiple independent C51 learners (Bellemare et al., 2017). This approach performs distributional Q-learning for each agent, and is able to deliver superior results. Lyu and Amato (2020) combined Hysteretic Q-Learning (HQL) (Matignon et al., 2007; Omidshafiei et al., 2017) with implicit quantile networks (IQN) (Dabney et al., 2018a) to separate the impacts of miscoordination among agents and the environmental stochasticity. Rowland et al. (2021) proposed a framework to investigate the difference between asymmetric update methods such as HQL, and distributional RL methods such as expectile distributional reinforcement learning (EDRL) (Rowland et al., 2019).

CTDE Value Function Factorization. Value function factorization methods aim to learn a factorization of the joint value function for a group of agents during centralized training. The learned factorization allows the policy of each agent to be executed in a decentralized fashion. The first proposed method in this category is VDN (Sunehag et al., 2018). VDN assumes the *Additivity* property, which constrains a joint action-value function to be the sum of multiple individual utilities. As a result, VDN employs a joint value network to accomplish such an objective. QMIX (Rashid et al., 2018) extends VDN to a larger set of tasks by assuming the *Monotonicity* property, which requires a joint action-value function

to be equivalent to at least one of the monotonic transformations of individual utilities. Based on this relaxed assumption, QMIX utilizes a monotonic mixing network to combine individual utilities in a monotonic fashion. QTRAN (Son et al., 2019) further transforms the joint action-value function into another easily factorizable one by introducing two additional regularization losses, and allows an even larger set of tasks to be factorized correctly. However, QTRAN is unstable during training (Rashid et al., 2018), which motivates several subsequent works dedicated to extending QMIX to solve the same set of tasks as QTRAN while maintaining the training stability. For example, W-QMIX (Rashid et al., 2020) treats QMIX as an operator, and projects non-monotonic tasks to QMIX’s problem domain. Qatten (Yang et al., 2020) introduces a multi-head attention mechanism to approximate a joint action-value function. QPLEX (Wang et al., 2021b) utilizes a duplex dueling mixing network for stabilizing QTRAN’s training process.

Other CTDE Methods. Aside from the general CTDE value function factorization methods, a number of methods have been proposed to introduce additional assumptions to further enhance the general CTDE methods. They include the use of intrinsic reward signals (Du et al., 2019), explicit communication channels (Zhang et al., 2019; Wang et al., 2020b), common knowledge shared among agents (de Witt et al., 2019; Wang et al., 2020a), prior knowledge about environments (Wang et al., 2020c), and so on. In addition, there has been a line of research that focuses on an actor-critic style training. For example, Multi-Agent Deep Deterministic Policy Gradient (MADDPG) (Lowe et al., 2017) uses a centralized critic to improve the value function approximation ability. The critic is augmented with extra information about all agents’ policies, while the actor only has access to its local information. Counterfactual Multi-Agent Policy Gradients (COMA) (Foerster et al., 2018) leverages a centralized critic that approximates a baseline for each agent. The baseline enables an agent to marginalize out its action while keeping the other agents’ actions fixed, allowing it to reason about counterfactual scenarios to facilitate its learning. Multi-Actor-Attention-Critic (MAAC) (Iqbal and Sha, 2019) constructs a critic for each agent with an attention mechanism, such that the critics are able to selectively pay attention to the information from the other agents. Moreover, Factored Multi-Agent Centralized Policy Gradients (FACMAC) (Peng et al., 2021) incorporates value function factorization methods into the centralized critic in MADDPG to further improve its performance. Please note that in this paper, we focus on improving value-based estimation for CTDE training. These methods are orthogonal to our work.

Distributional Value Function Factorization. Similar to our work, QR-MIX (Hu et al., 2020) and RMIX (Qiu et al., 2021) also utilized distributional RL to increase their performance in SMAC. QR-MIX is a QMIX variant which models the mixing network as an IQN. Instead of decomposing the joint distribution into individual utility distributions, it treats the utilities as Dirac distributions. On the other hand, RMIX is also a QMIX variant, but replaces the inputs and outputs of the mixing network with the Conditional Value at Risk (CVaR) measures. The major difference between our work and the above two methods is the ability to factorize the joint value distribution into individual utility distributions, while these two methods are not capable of factorizing distributions.

3. Background

In this section, we introduce the essential background material for understanding the contents of this paper. We first define the problem formulation of cooperative MARL and CTDE. Next, we describe the conventional formulation of IGM and the value function factorization methods. Then, we walk through the concepts of distributional RL, quantile function, as well as quantile regression, which are the fundamental concepts frequently mentioned in this paper. After that, we explain IQN, a key approach adopted in this paper for approximating quantiles. Finally, we bring out the concept of quantile mixture, which is leveraged by DFAC for factorizing return distributions.

3.1 Cooperative MARL and CTDE

In this work, we consider a fully cooperative MARL environment modeled as a decentralized and partially observable Markov Decision Process (Dec-POMDP) (Oliehoek and Amato, 2016) with stochastic rewards, which is described as a tuple $\langle \mathbb{S}, \mathbb{K}, \mathbb{O}_{\text{jt}}, \mathbb{U}_{\text{jt}}, P, O, R, \gamma \rangle$ and is defined as follows:

- \mathbb{S} is the finite set of global states in the environment, where $s' \in \mathbb{S}$ denotes the next state of the current state $s \in \mathbb{S}$. The state information is optionally available during training, but is not available to the agents during execution.
- $\mathbb{K} = \{1, \dots, K\}$ is the set of K agents. We use $k \in \mathbb{K}$ to denote the index of the agent.
- $\mathbb{O}_{\text{jt}} = \prod_{k \in \mathbb{K}} \mathbb{O}_k$ is the set of joint observations. At each timestep, a joint observation $\mathbf{o} = \langle o_1, \dots, o_K \rangle \in \mathbb{O}_{\text{jt}}$ is provided by the environment to the K agents. Each agent k is only able to observe its own individual observation $o_k \in \mathbb{O}_k$.
- $\mathbb{H}_{\text{jt}} = \prod_{k \in \mathbb{K}} \mathbb{H}_k$ is the set of joint action-observation histories. The joint history $\mathbf{h} = \langle h_1, \dots, h_K \rangle \in \mathbb{H}_{\text{jt}}$ concatenates all the perceived observations and the performed actions before a certain timestep, where $h_k \in \mathbb{H}_k$ represents the action-observation history from agent k .
- $\mathbb{U}_{\text{jt}} = \prod_{k \in \mathbb{K}} \mathbb{U}_k$ is the set of joint actions. At each timestep, the entire group of the agents take a joint action \mathbf{u} , where $\mathbf{u} = \langle u_1, \dots, u_K \rangle \in \mathbb{U}_{\text{jt}}$. The individual action $u_k \in \mathbb{U}_k$ of each agent k is determined based on its stochastic policy $\pi_k(u_k|h_k) : \mathbb{H}_k \times \mathbb{U}_k \rightarrow [0, 1]$, expressed as $u_k \sim \pi_k(\cdot|h_k)$. Similarly, in single agent scenarios, we use u and u' to denote the actions of the agent at state s and s' under its policy π , respectively.
- $\mathbb{T} = \{1, \dots, T\}$ represents the set of timesteps with horizon T , where the index of the current timestep is denoted as $t \in \mathbb{T}$. Please note that s^t , \mathbf{o}^t , \mathbf{h}^t , and \mathbf{u}^t correspond to the state, joint observation, joint action-observation history, and joint action at timestep t , respectively.
- The transition function $P(s'|s, \mathbf{u}) : \mathbb{S} \times \mathbb{U}_{\text{jt}} \times \mathbb{S} \rightarrow [0, 1]$ specifies the state transition probabilities. Given s and \mathbf{u} , the next state is represented as $s' \sim P(\cdot|s, \mathbf{u})$.
- The observation function $O(\mathbf{o}|s) : \mathbb{O}_{\text{jt}} \times \mathbb{S} \rightarrow [0, 1]$ specifies the joint observation probabilities. Given s , the joint observation is represented as $\mathbf{o} \sim O(\cdot|s)$.
- $R(r|s, \mathbf{u}) : \mathbb{S} \times \mathbb{U}_{\text{jt}} \times \mathbb{R} \rightarrow [0, 1]$ is the joint reward function shared among all the agents. Given s , the team reward is expressed as $r \sim R(\cdot|s, \mathbf{u})$.
- $\gamma \in \mathbb{R}$ is the discount factor with its value within $[0, 1)$.

Under such an MARL formulation, this work concentrates on CTDE-based value function factorization methods, where the agents are trained in a centralized fashion and executed in a decentralized manner. In other words, the joint action-observation history \mathbf{h} is available during the learning processes of individual policies $[\pi_k]_{k \in \mathbb{K}}$. During execution, each agent’s policy π_k only conditions on its observation history h_k .

3.2 IGM and Factorizable Tasks

IGM is necessary for value function factorization (Son et al., 2019). For a joint action-value function $Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) : \mathbb{H}_{\text{jt}} \times \mathbb{U}_{\text{jt}} \rightarrow \mathbb{R}$, if there exist K separate utility functions $[Q_k(h_k, u_k) : \mathbb{H}_k \times \mathbb{U}_k \rightarrow \mathbb{R}]_{k \in \mathbb{K}}$ such that the following condition holds:

$$\arg \max_{\mathbf{u}} Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) = \begin{pmatrix} \arg \max_{u_1} Q_1(h_1, u_1) \\ \vdots \\ \arg \max_{u_K} Q_K(h_K, u_K) \end{pmatrix}, \quad (1)$$

then $[Q_k]_{k \in \mathbb{K}}$ are said to satisfy IGM for Q_{jt} under \mathbf{h} . Under this condition, we also say that $Q_{\text{jt}}(\mathbf{h}, \mathbf{u})$ is factorized by $[Q_k(h_k, u_k)]_{k \in \mathbb{K}}$ (Son et al., 2019). If Q_{jt} in a given task is factorizable under all $\mathbf{h} \in \mathbb{H}_{\text{jt}}$, we say that the task is factorizable. Intuitively, factorizable tasks indicate that there exists a factorization such that each agent can select the greedy action according to their individual utilities $[Q_k]_{k \in \mathbb{K}}$ independently in a decentralized fashion. This enables the optimal individual actions to implicitly achieve the optimal joint action across the K agents. Since there is no individual reward, the factorized utilities do not estimate expected returns on their own (Guestrin et al., 2001) and are different from the value function definition commonly used in SARL.

3.3 Advantage-based IGM

In addition to Q-value based IGM described in Section 3.2, there exist a variant, called *advantage-based IGM* (Wang et al., 2021b), which is defined as the following:

$$\arg \max_{\mathbf{u}} A_{\text{jt}}(\mathbf{h}, \mathbf{u}) = \begin{pmatrix} \arg \max_{u_1} A_1(h_1, u_1) \\ \vdots \\ \arg \max_{u_K} A_K(h_K, u_K) \end{pmatrix}, \quad (2)$$

where the advantage functions $[A_k(h_k, u_k)]_{k \in \mathbb{K}}$ can be derived from $[Q_k(h_k, u_k)]_{k \in \mathbb{K}}$ and the state-value functions $[V_k(h_k)]_{k \in \mathbb{K}}$ as:

$$A_k(h_k, u_k) = Q_k(h_k, u_k) - V_k(h_k), \forall k \in \mathbb{K}, \quad (3)$$

$$V_k(h_k) = \max_{u_k} Q_k(h_k, u_k), \forall k \in \mathbb{K}. \quad (4)$$

Eq. (2) has been proved to be equivalent to Eq. (1) (Wang et al., 2021b), since the argmax operations over the actions do not depend on $[V_k(h_k)]_{k \in \mathbb{K}}$, but only on $[A_k(h_k, u_k)]_{k \in \mathbb{K}}$.

3.4 Value Function Factorization Methods

Based on IGM, value function factorization methods enable centralized training for factorizable tasks, while maintaining the ability for decentralized execution. In this work, we consider three such methods: VDN, QMIX, and QPLEX, where their expressive power (or expressiveness) can be sorted as: QPLEX > QMIX > VDN. A factorization function being more expressive means that it is able to solve a larger subset of factorizable tasks. VDN solves a subset of factorizable tasks that satisfies *Additivity* by defining the factorization function as the sum of individual utilities:

$$Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) = \sum_{k=1}^K Q_k(h_k, u_k). \quad (5)$$

QMIX solves a larger subset of factorizable tasks that satisfies *Monotonicity* by defining the factorization function as a monotonic combination of the individual utilities:

$$Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) = M(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s), \quad (6)$$

where M is a learnable monotonic mixing network that satisfies $\frac{\partial M}{\partial Q_k} \geq 0, \forall k \in \mathbb{K}$, and can condition on the state s if the information is available during training. QPLEX solves all factorizable tasks (Wang et al., 2021b) by defining the factorization function as follows:

$$Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) = \sum_{k=1}^K V_k(\mathbf{h}) + \sum_{k=1}^K \lambda_k(\mathbf{h}, \mathbf{u}) A_k(\mathbf{h}, u_k). \quad (7)$$

where $V_k(\mathbf{h}) = M_k(V_k(h_k)|s)$, $A_k(\mathbf{h}, u_k) = M_k(A_k(h_k, u_k)|s)$, and $\lambda_k(\mathbf{h}, \mathbf{u}) > 0$. In addition to the learnable monotonic networks $[M_k]_{k \in \mathbb{K}}$, such a factorization further learns a positive weight λ_k for each individual advantage. This flexibility enables QPLEX to correct the discrepancy between the joint Q-value and the individual utilities when a task does not satisfy *Additivity*. All of these three factorization functions satisfy the IGM condition (Son et al., 2019), which indicates that the individual utilities $[Q_k^{(t)}]_{k \in \mathbb{K}}$ satisfy IGM for $Q_{\text{jt}}^{(t)}$ under all \mathbf{h} throughout the entire training process (i.e., for all t). The joint Q-value Q_{jt} is learned iteratively by updating the joint network with the Bellman optimality operator \mathcal{T}^* based on the joint reward signal:

$$Q_{\text{jt}}^{(i+1)}(\mathbf{h}, \mathbf{u}) \leftarrow \mathcal{T}^* Q_{\text{jt}}^{(i)}(\mathbf{h}, \mathbf{u}), \quad (8)$$

where i denotes the index of iterations, $\mathcal{T}^* Q_{\text{jt}}^{(i)}(\mathbf{h}, \mathbf{u}) = \mathbb{E}[R(s, \mathbf{u})] + \gamma Q_{\text{jt}}^{(i)}(s', \mathbf{u}^*)$, and $\mathbf{u}^* = \arg \max_{\mathbf{u}'} Q_{\text{jt}}^{(i)}(s', \mathbf{u}')$ is the optimal action at state s' . Eq. (8) enables the joint network to be trained in the same fashion as the single-agent DQN (Mnih et al., 2015), allowing Q_{jt} to approximate the optimal joint Q-value function Q_{jt}^* while satisfying the IGM condition. This update process of minimizing the Temporal Difference (TD) error implicitly realizes an effective counterfactual credit assignment by factorizing Q_{jt} into individual utilities $[Q_k]_{k \in \mathbb{K}}$ (Wang et al., 2021a). The factorized utilities are then utilized to determine each agent’s policy during execution.

The representational capacity of the joint network is determined by two factors: (1) the network’s (or the model’s) capacity, which is limited by the number of learnable parameters,

and (2) the expressiveness of a factorization function, which is dependent on the selected factorization method. For example, QPLEX has been proved to be more expressive than QMIX (Wang et al., 2021b). Due to the theoretical limitation of QMIX’s expressiveness, there exist certain tasks that QMIX are unable to solve by simply minimizing the TD errors of the joint Q-values, even if QMIX is equipped with unlimited model capacity.

3.5 Distributional RL

For notational simplicity, we consider a degenerated case with only a single agent, and the environment is fully observable until the end of Section 3.8. Distributional RL generalizes classic expected RL methods by capturing the full return distribution $Z(s, u)$ instead of the expected return $Q(s, u)$, allowing an agent to learn better state representations (Bellemare et al., 2022) that lead to better performance in various single-agent domains (Bellemare et al., 2017, 2019; Dabney et al., 2018b,a; Rowland et al., 2019; Yang et al., 2019; Nguyen et al., 2021). Modeling the full return distributions can be viewed as an auxiliary task that benefits representation learning during training, where the information other than the expected returns can be discarded afterward. In addition, distributional RL enables applications that require the information of the full return distributions, such as exploration (Nikolov et al., 2019; Zhang and Yao, 2019; Mavrin et al., 2019) and risk-aware policies (Xia, 2020). Distributional RL defines the distributional Bellman operator \mathcal{T}^π as follows:

$$\mathcal{T}^\pi Z(s, u) \stackrel{D}{=} R(s, u) + \gamma Z(s', u'), \quad (9)$$

and the distributional Bellman optimality operator \mathcal{T}^* as:

$$\mathcal{T}^* Z(s, u) \stackrel{D}{=} R(s, u) + \gamma Z(s', u'^*), \quad (10)$$

where $u'^* = \arg \max_{u'} \mathbb{E}[Z(s', u')]$ is the optimal action at state s' , and the expression $X \stackrel{D}{=} Y$ denotes that random variables X and Y follow the same distribution. Given some initial distribution Z_0 , applying different Bellman operators lead to different results. If \mathcal{T}^π is applied repeatedly, Z converges to Z^π in p -Wasserstein distance for all $p \in [1, \infty)$ under π . On the other hand, if \mathcal{T}^* is applied instead, Z alternates among the optimal return distributions in the set $\mathcal{Z}^* := \{Z^{\pi^*} : \pi^* \in \Pi^*\}$, where Π^* denotes the set of optimal policies (Bellemare et al., 2017). The p -Wasserstein distance between the probability distributions of random variables X, Y is defined as:

$$W_p(X, Y) = \left(\int_0^1 |F_X^{-1}(\omega) - F_Y^{-1}(\omega)|^p d\omega \right)^{1/p}, \quad (11)$$

where (F_X^{-1}, F_Y^{-1}) are quantile functions of (X, Y) .

3.6 Categorical Distribution and Heuristic Projection

The distributional Bellman optimality operator \mathcal{T}^* lays the foundation for C51 (Bellemare et al., 2017), the first distributional RL algorithm. It models the return distribution as a categorical distribution, where the number of atoms n (i.e., categories) within the range $[V_{\text{MIN}}, V_{\text{MAX}}]$ and the distance between them $\Delta z = \frac{V_{\text{MAX}} - V_{\text{MIN}}}{n-1}$ are hyperparameters. Since

C51 can only approximate the return distributions with a finite support (i.e., the set of atoms: $\{z_i = V_{\text{MIN}} + i\Delta z : 0 \leq i < n\}$), a heuristic projection is required to project the Bellman target distribution onto the pre-defined support. Specifically, C51 aims to minimize the KL divergence:

$$\mathcal{L}(s, u, r, s') = D_{\text{KL}}(\text{proj}(\mathcal{T}^*Z(s, u)) || Z(s, u)), \quad (12)$$

where the heuristic projection function $\text{proj}(\cdot)$ approximates $\mathcal{T}^*Z(s, u)$ by distributing the probability of each atoms to their immediate neighbors on the support of C51:

$$[\text{proj}(\mathcal{T}^*Z(s, u))]_i = \sum_{j=0}^{n-1} \left[1 - \frac{|\llbracket \mathcal{T}^*z_j \rrbracket_{V_{\text{MIN}}}^{V_{\text{MAX}}} - z_i|}{\Delta z} \right]_0^1 p_j(s', u^*), \quad (13)$$

where $[\cdot]_a^b$ is a clipping function that bounds its argument in $[a, b]$, and $p_j(s', u^*)$ is the probability mass of the j^{th} atom. Unfortunately, the heuristic projection used in Eq. (13) introduces variance (Rowland et al., 2019) and does not guarantee convergence (Dabney et al., 2018b). In addition, C51 usually requires a large number of categories, and often necessitates re-adjusting the support range for tasks with different reward magnitudes, making it inappropriate for practical usages. Therefore, its follow-up works turn to approximate return distributions with quantile functions instead of categorical distributions.

3.7 Quantile Function and Quantile Regression

The relationship between the cumulative distribution function (CDF) F_X and the quantile function F_X^{-1} (i.e., the generalized inverse CDF) of a random variable X is formulated as:

$$F_X^{-1}(\omega) = \inf\{x \in \mathbb{R} : \omega \leq F_X(x)\}, \forall \omega \in [0, 1], \quad (14)$$

where ω represents the quantile. The expectation of X expressed in terms of $F_X^{-1}(\omega)$ is:

$$\mathbb{E}[X] = \int_0^1 F_X^{-1}(\omega) \, d\omega. \quad (15)$$

Dabney et al. (2018a) model the value function as a quantile function $F^{-1}(s, u|\omega)$, and use a pair-wise sampled temporal difference (TD) error δ for two quantile samples $\omega, \omega' \sim U([0, 1])$ to optimize the value function. The TD error δ is defined as:

$$\delta^{\omega, \omega'} = r + \gamma F^{-1}(s', u^*|\omega') - F^{-1}(s, u|\omega). \quad (16)$$

Based on Eq. (16), the pair-wise loss ρ_ω^κ is formulated based on the Huber quantile regression loss \mathcal{L}_κ (Dabney et al., 2018b) with threshold $\kappa = 1$, and is expressed as follows:

$$\rho_\omega^\kappa(\delta^{\omega, \omega'}) = |\omega - \mathbb{I}\{\delta^{\omega, \omega'} < 0\}| \frac{\mathcal{L}_\kappa(\delta^{\omega, \omega'})}{\kappa}, \text{ where} \quad (17)$$

$$\mathcal{L}_\kappa(\delta^{\omega, \omega'}) = \begin{cases} \frac{1}{2}(\delta^{\omega, \omega'})^2, & \text{if } |\delta^{\omega, \omega'}| \leq \kappa \\ \kappa(|\delta^{\omega, \omega'}| - \frac{1}{2}\kappa), & \text{otherwise} \end{cases}. \quad (18)$$

Given N quantile samples $[\omega_i]_{i=1}^N$ to be optimized with regard to N' target quantile samples $[\omega_j]_{j=1}^{N'}$, the total loss $\mathcal{L}(s, u, r, s')$ is defined as the sum of the pair-wise losses, and is expressed as the following:

$$\mathcal{L}(s, u, r, s') = \frac{1}{N'} \sum_{i=1}^N \sum_{j=1}^{N'} \rho_{\omega_i}^{\kappa}(\delta^{\omega_i, \omega'_j}). \quad (19)$$

3.8 Implicit Quantile Network

In order to realize the concepts discussed in the previous section, implicit quantile network (IQN) (Dabney et al., 2018a) proposes to approximate the return distribution $Z(s, u)$ by an implicit quantile function $F^{-1}(s, u|\omega) = g(\psi(s), \phi(\omega))_u$ for functions g , ψ , and ϕ . The subscript u denotes a certain action, which corresponds to the u^{th} element in the output vector of g . IQN is a type of universal value function approximator (UVFA) (Schaul et al., 2015). It employs a light-weight architecture which generalizes its predictions across states $s \in \mathbb{S}$ and goals $\omega \in [0, 1]$, with the goals defined as different quantiles of the return distribution. To mitigate *spectral bias* (Rahaman et al., 2019), ϕ first maps the input scalar ω to an m -dimensional vector by a high frequency function $[\cos(\pi i \omega)]_{i=0}^{m-1}$, followed by a single hidden layer with weights $[w_{ij}]$ and biases $[b_j]$ to produce a quantile embedding $\phi(\omega) = [\phi(\omega)_j]_{j=0}^{\dim(\phi(\omega))-1}$. The expression of $\phi(\omega)_j$ can be represented as the following:

$$\phi(\omega)_j := \text{ReLU}\left(\sum_{i=0}^{m-1} \cos(\pi i \omega) w_{ij} + b_j\right), \quad (20)$$

where $m = 64$. Then, $\phi(\omega)$ is combined with the state embedding $\psi(s)$ by the element-wise (Hadamard) product (\odot), expressed as $g := \psi \odot \phi$, where $\dim(\phi(\omega)) = \dim(\psi(s))$. The loss of IQN is defined as Eq. (19) by sampling a batch of N and N' quantiles from the policy network and the target network respectively. During execution, the action with the largest expected return $Q(s, u)$ is chosen. Since IQN does not model the expected return explicitly, $Q(s, u)$ is approximated by calculating the mean of the sampled returns through \hat{N} quantile samples $\hat{\omega}_i \sim U([0, 1])$, $\forall i \in [1, \hat{N}]$ based on Eq. (15), expressed as follows:

$$Q(s, u) = \int_0^1 F^{-1}(s, u|\omega) d\omega \approx \frac{1}{\hat{N}} \sum_{i=1}^{\hat{N}} F^{-1}(s, u|\hat{\omega}_i). \quad (21)$$

3.9 Quantile Mixture

Multiple quantile functions (e.g., IQNs) sharing the same quantile ω may be combined into a single quantile function $F^{-1}(\omega)$, in a form of quantile mixture expressed as follows:

$$F^{-1}(\omega) = \sum_{k=1}^K \beta_k F_k^{-1}(\omega), \quad (22)$$

where $[F_k^{-1}(\omega)]_{k \in \mathbb{K}}$ are quantile functions, and $[\beta_k]_{k \in \mathbb{K}}$ are model parameters (Karvanen, 2006). The condition for $[\beta_k]_{k \in \mathbb{K}}$ is that $F^{-1}(\omega)$ must satisfy the properties of a quantile

function. The concept of quantile mixture must not be confused with the mixture of multiple probability density functions (PDFs):

$$f(x) = \sum_{k=1}^K \alpha_k f_k(x), \tag{23}$$

where $[f_k(x)]_{k \in \mathbb{K}}$ are PDFs, $\sum_{k=1}^K \alpha_k = 1$, and $\alpha_k \geq 0$. Summing PDFs is equivalent to making a random choice from a collection of random variables, whereas adding quantile functions combines the random variables with a specific dependence structure (Eq. (22)).

4. Methodology

In this section, we walk through the proposed DFAC framework and its derivation procedure. We first discuss a naive distributional factorization approach and its limitation in Section 4.1. Then, we introduce the DFAC framework to address the limitation, and show that DFAC is able to generalize distributional RL to all factorizable tasks in Section 4.2. After that, several practical implementations of DFAC are presented in Section 4.3. Finally, DDN, DMIX, and DPLEX, which are the DFAC variants of VDN, QMIX, and QPLEX, respectively, are discussed in Section 4.4.

4.1 Distributional IGM

Since the IGM condition is required for value function factorization to work in fully cooperative tasks, a distributional factorization that satisfies IGM is essential for factorizing return distributions. We first discuss a naive distributional factorization approach that simply replaces deterministic utilities Q with stochastic utilities Z . Then, we provide a theorem to show that the naive approach is insufficient to guarantee the IGM condition for all factorization functions in general.

Definition 1 (Distributional IGM) *A finite number of individual stochastic utilities $[Z_k(h_k, u_k)]_{k \in \mathbb{K}}$ are said to satisfy Distributional IGM (DIGM) for a stochastic joint action-value function $Z_{\text{jt}}(\mathbf{h}, \mathbf{u})$ under \mathbf{h} , if their expectations $[\mathbb{E}[Z_k(h_k, u_k)]]_{k \in \mathbb{K}}$ satisfy IGM for $\mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})]$ under \mathbf{h} , represented as follows:*

$$\arg \max_{\mathbf{u}} \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] = \begin{pmatrix} \arg \max_{u_1} \mathbb{E}[Z_1(h_1, u_1)] \\ \vdots \\ \arg \max_{u_K} \mathbb{E}[Z_K(h_K, u_K)] \end{pmatrix}.$$

By replacing Q in Eq. (1) with $\mathbb{E}[Z]$, IGM can be extended to Distributional IGM. The Advantage-based IGM can also be extended to its distributional version in a similar way. In the following, we use Q and $\mathbb{E}[Z]$ interchangeably, specifically, $Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) = \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})]$, and $Q_k(h_k, u_k) = \mathbb{E}[Z_k(h_k, u_k)]$, $\forall k \in \mathbb{K}$.

Proposition 1 *DIGM is equivalent to IGM (Eq. (1)).*

Proof

$$\begin{aligned} \arg \max_{\mathbf{u}} \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] &= \begin{pmatrix} \arg \max_{u_1} \mathbb{E}[Z_1(h_1, u_1)] \\ \vdots \\ \arg \max_{u_K} \mathbb{E}[Z_K(h_K, u_K)] \end{pmatrix} \\ \Leftrightarrow \arg \max_{\mathbf{u}} Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) &= \begin{pmatrix} \arg \max_{u_1} Q_1(h_1, u_1) \\ \vdots \\ \arg \max_{u_K} Q_K(h_K, u_K) \end{pmatrix}. \end{aligned}$$

■

The Advantage-based IGM is also equivalent to its distributional version, which can be proved in the same manner. Following the above definition of DIGM, we next discuss the reason why simply replacing Q with Z (i.e., the naive approach) is insufficient to guarantee the DIGM condition by considering *Additivity* (in Eq. (5)) and *Monotonicity* (in Eq. (6)) in two propositions.

Proposition 2 *Additivity for utility distributions:*

$$Z_{\text{jt}}(\mathbf{h}, \mathbf{u}) \stackrel{D}{=} \Psi(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)) \stackrel{D}{=} \sum_{k \in \mathbb{K}} Z_k(h_k, u_k)$$

is a sufficient condition for DIGM.

Proof By linearity of expectation, the following equation holds:

$$\begin{aligned} \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] &= \mathbb{E}\left[\sum_{k \in \mathbb{K}} Z_k(h_k, u_k)\right] = \sum_{k \in \mathbb{K}} \mathbb{E}[Z_k(h_k, u_k)] \\ \Leftrightarrow Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) &= \sum_{k \in \mathbb{K}} Q_k(h_k, u_k). \end{aligned}$$

Since *Additivity* is a sufficient condition for IGM:

$$\begin{aligned} \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] &= \Psi(\mathbb{E}[Z_1(h_1, u_1)], \dots, \mathbb{E}[Z_K(h_K, u_K)]) = \sum_{k \in \mathbb{K}} \mathbb{E}[Z_k(h_k, u_k)] \\ \Leftrightarrow Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) &= \Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)) = \sum_{k \in \mathbb{K}} Q_k(h_k, u_k), \end{aligned}$$

$[\mathbb{E}[Z_k(h_k, u_k)]]_{k \in \mathbb{K}}$ satisfy IGM for $\mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})]$ under \mathbf{h} . According to Definition 1 and Proposition 1, $[Z_k(h_k, u_k)]_{k \in \mathbb{K}}$ satisfy DIGM for $Z_{\text{jt}}(\mathbf{h}, \mathbf{u})$ under \mathbf{h} . As a result, *Additivity* for utility distributions is a sufficient condition for DIGM. ■

Simply replacing Q with Z (i.e., the naive approach) satisfies DIGM for *Additivity* based on linearity of expectations. However, linearity of expectations does not hold in general for monotonic transformations:

$$\begin{aligned} Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) &= \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] \\ &= \mathbb{E}[M(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s)], \forall M \\ &\neq M(\mathbb{E}[Z_1(h_1, u_1)], \dots, \mathbb{E}[Z_K(h_K, u_K)]|s), \forall M \\ &= M(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s), \forall M. \end{aligned}$$

Next, we provide a counterexample to demonstrate this property.

Proposition 3 *Monotonicity for utility distributions:*

$$\begin{aligned} Z_{\text{jt}}(\mathbf{h}, \mathbf{u}) &\stackrel{D}{=} \Psi(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s) \\ &\stackrel{D}{=} M(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s), \end{aligned}$$

where M is a monotonic transformation that satisfies $\frac{\partial M}{\partial Q_k} \geq 0, \forall k \in \mathbb{K}$, is not a sufficient condition for DIGM, even though the condition may satisfy DIGM for special cases of M and $[Z_k(h_k, u_k)]_{k \in \mathbb{K}}$.

Proof We consider a degenerated case and prove the theorem by contradiction. Consider a case where there is only a single agent ($K = 1$), with a single fully observable state and an exponential transformation $M(Z_1(h_1, u_1)|s) = \exp(Z_1(h_1, u_1))$. The (joint) action space of this case consists of two (joint) actions: $\mathbb{U}_{\text{jt}} = \mathbb{U}_1 = \{u_1^*, u_1'\}$, where u_1^* is the optimal action (with expected return 2) and u_1' is the suboptimal action (with expected return 1.5). Given the probability mass function (PMF) of $Z_1(h_1, u_1^*)$ as:

$$p_1(z) = \begin{cases} 1 & \text{if } z = 2 \\ 0 & \text{otherwise,} \end{cases}$$

and the PMF of $Z_1(h_1, u_1')$ as:

$$p_2(z) = \begin{cases} 0.5 & \text{if } z = 0 \\ 0.5 & \text{if } z = 3 \\ 0 & \text{otherwise.} \end{cases}$$

Based on this special case, we calculate the optimal actions before applying and after applying the monotonic transformation as follows:

$$\begin{cases} \mathbb{E}[Z_1(h_1, u_1^*)] = 1 \cdot 2 = 2 \\ \mathbb{E}[Z_1(h_1, u_1')] = 0.5 \cdot 0 + 0.5 \cdot 3 = 1.5 \\ \arg \max_{u_1} \mathbb{E}[Z_1(h_1, u_1)] = u_1^* \\ \mathbb{E}[M(Z_1(h_1, u_1^*)|s)] = \mathbb{E}[\exp(Z_1(h_1, u_1^*))] = e^2 \approx 7.39 \\ \mathbb{E}[M(Z_1(h_1, u_1')|s)] = \mathbb{E}[\exp(Z_1(h_1, u_1'))] = 0.5 \cdot e^0 + 0.5 \cdot e^3 \approx 10.54 \\ \arg \max_{u_1} \mathbb{E}[M(Z_1(h_1, u_1)|s)] = u_1' \end{cases}$$

Assume, to the contrary, that *Monotonicity* for utility distributions is a sufficient condition for DIGM. By the definition of DIGM (Definition 1):

$$\begin{aligned} \arg \max_{\mathbf{u}} \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] &= (\arg \max_{u_1} \mathbb{E}[Z_1(h_1, u_1)]) \\ \Rightarrow \arg \max_{u_1} \mathbb{E}[M(Z_1(h_1, u_1)|s)] &= \arg \max_{u_1} \mathbb{E}[Z_1(h_1, u_1)] \\ \Rightarrow u_1' = u_1^* & (\Rightarrow \text{contradiction}). \end{aligned}$$

A contradiction occurs since $u_1' \neq u_1^*$, showing that *Monotonicity* is not a sufficient condition for DIGM. Since there exists a case where DIGM does not hold for $K = 1$, it certainly

does not hold for all $K \in \mathbb{Z}^+$. ■

The reason that the naive approach does not hold for *Monotonicity* is because the monotonic function (i.e., \exp) in the above example reshapes the distributions in a non-linear manner. This may potentially cause reordering of the actions' expected returns. For example, consider a case of two distributions, a left-skewed and a right-skewed distributions with the same expectation. After applying the exponential transformation, the expectation of the left-skewed distribution becomes larger than that of the right-skewed distribution. Such an effect only exists in non-linear transformations and does not exist in linear transformations.

Theorem 1 *Given a deterministic joint action-value function Q_{jt} , a stochastic joint action-value function Z_{jt} , and a factorization function Ψ for deterministic utilities:*

$$Q_{jt}(\mathbf{h}, \mathbf{u}) = \Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s),$$

such that $[Q_k]_{k \in \mathbb{K}}$ satisfy IGM for Q_{jt} under \mathbf{h} , the following distributional factorization:

$$Z_{jt}(\mathbf{h}, \mathbf{u}) \stackrel{D}{=} \Psi(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s).$$

is insufficient to guarantee that $[Z_k]_{k \in \mathbb{K}}$ satisfy DIGM for Z_{jt} under \mathbf{h} .

Proof A contradiction is provided by Proposition 3. ■

According to the observations in Theorem 1, an alternative strategy (other than the naive approach) for modifying the factorization functions is necessary to meet the requirements of satisfying DIGM for stochastic joint distributions.

4.2 Mean-Shape Decomposition and the DFAC Framework

We propose Mean-Shape Decomposition and the DFAC framework to ensure that DIGM is satisfied for utility distributions. We first define Mean-Shape Decomposition as follows:

Definition 2 (Mean-Shape Decomposition) *For any given random variable Z , there exists a unique decomposition defined as follows:*

$$\begin{aligned} Z &= \mathbb{E}[Z] + (Z - \mathbb{E}[Z]) \\ &= Z_{\text{mean}} + Z_{\text{shape}}, \end{aligned}$$

where $\text{Var}(Z_{\text{mean}}) = 0$ and $\mathbb{E}[Z_{\text{shape}}] = 0$. In this paper, this is called the *Mean-Shape Decomposition* of Z .

Based on Mean-Shape Decomposition, we propose DFAC to decompose a joint return distribution Z_{jt} into its deterministic part Z_{mean} (i.e., the expected value) and stochastic part Z_{shape} (i.e., the higher moments). The two components Z_{mean} and Z_{shape} are approximated by two different functions Ψ and Φ , respectively. The factorization function Ψ is responsible for factorizing the expectation of Z_{jt} , while the shape function Φ is utilized to factorize the shape of Z_{jt} . Since the main objective of modeling the return distribution is to assist non-linear approximation of the expectation of Z_{jt} (Lyle et al., 2019; Bellemare et al., 2022) rather than accurately model the shape of Z_{jt} , Φ is allowed to roughly factorize the shape of Z_{jt} .

Theorem 2 (DFAC Theorem) Consider a deterministic joint action-value function Q_{jt} , a stochastic joint action-value function Z_{jt} , and a factorization function Ψ for deterministic utilities:

$$Q_{\text{jt}}(\mathbf{h}, \mathbf{u}) = \Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s),$$

such that $[Q_k]_{k \in \mathbb{K}}$ satisfy IGM for Q_{jt} under \mathbf{h} . By Mean-Shape Decomposition, the following distributional factorization:

$$\begin{aligned} Z_{\text{jt}}(\mathbf{h}, \mathbf{u}) &= \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] + (Z_{\text{jt}}(\mathbf{h}, \mathbf{u}) - \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})]) \\ &= Z_{\text{mean}}(\mathbf{h}, \mathbf{u}) + Z_{\text{shape}}(\mathbf{h}, \mathbf{u}) \\ &\stackrel{D}{=} \Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s) + \Phi(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s), \end{aligned}$$

is sufficient to guarantee that $[Z_k]_{k \in \mathbb{K}}$ satisfy DIGM for Z_{jt} under \mathbf{h} , where $\mathbb{E}[\Phi(Z_1, \dots, Z_K|s)] = 0$ for all $[Z_k]_{k \in \mathbb{K}}$.

Proof Based on Mean-Shape decomposition:

$$\begin{aligned} &\arg \max_{\mathbf{u}} \{\mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})]\} \\ &= \arg \max_{\mathbf{u}} \{\mathbb{E}[Z_{\text{mean}}(\mathbf{h}, \mathbf{u}) + Z_{\text{shape}}(\mathbf{h}, \mathbf{u})]\} \\ &= \arg \max_{\mathbf{u}} \{\mathbb{E}[Z_{\text{mean}}(\mathbf{h}, \mathbf{u})] + \mathbb{E}[Z_{\text{shape}}(\mathbf{h}, \mathbf{u})]\} \\ &= \arg \max_{\mathbf{u}} \{\mathbb{E}[\Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s)] + \mathbb{E}[\Phi(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s)]\} \\ &= \arg \max_{\mathbf{u}} \{\Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s) + 0\} \\ &= \arg \max_{\mathbf{u}} \{\Psi(Q_1(h_1, u_1), \dots, Q_K(h_K, u_K)|s)\} \\ &= \begin{pmatrix} \arg \max_{u_1} Q_1(h_1, u_1) \\ \vdots \\ \arg \max_{u_K} Q_K(h_K, u_K) \end{pmatrix} \\ &\Rightarrow \arg \max_{\mathbf{u}} \mathbb{E}[Z_{\text{jt}}(\mathbf{h}, \mathbf{u})] = \begin{pmatrix} \arg \max_{u_1} \mathbb{E}[Z_1(h_1, u_1)] \\ \vdots \\ \arg \max_{u_K} \mathbb{E}[Z_K(h_K, u_K)] \end{pmatrix}. \end{aligned}$$

The above derivation demonstrates that $[Z_k]_{k \in \mathbb{K}}$ satisfy DIGM for Z_{jt} under \mathbf{h} . ■

Theorem 2 reveals that the choice of Ψ determines whether DIGM holds, regardless of the choice of Φ , as long as $\mathbb{E}[\Phi(Z_1, \dots, Z_K|s)] = 0$ for all $[Z_k]_{k \in \mathbb{K}}$. Under this setting, any factorization function of deterministic variables can be extended to a factorization function of stochastic variables. Such a decomposition enables approximation of joint distributions for all factorizable tasks under appropriate choices of Ψ and Φ . The methods extended by DFAC is called the DFAC variants, which have the same expressiveness as the original unextended versions.

4.3 Practical Implementation Choices of the Shape Function

In this section, we discuss the implementation choices of the shape function based on two representative distributional algorithms: C51 and IQN. Theoretically, the shape function

can be arbitrarily complex as long as the output has a zero mean. In practice, however, such a complicated network is unnecessary, since learning the shape of distribution is only an auxiliary task to facilitate representation learning, as mentioned in Section 3.5. For the sake of simplicity, we choose the shape function to be additive and non-learnable: $Z_{\text{shape}} = \Phi(Z_1(h_1, u_1), \dots, Z_K(h_K, u_K)|s) = \sum_{k \in \mathbb{K}} (Z_k - \mathbb{E}[Z_k])$.

4.3.1 SHAPE FUNCTION FOR C51

In this section, we demonstrate the implementation of the shape function for C51, which models the utilities as PMFs. We denote the PMFs of the utilities as $[f_k]_{k \in \mathbb{K}}$, and the PMF of the joint return as f . For the sake of simplicity, we first consider the case for two agents ($K = 2$). To exactly factorize the total return ($Z = Z_1 + Z_2$), the PMF f of the joint return Z can be expressed as:

$$\begin{aligned} f(x) &= \Pr(Z = x) \\ &= \Pr(Z_1 + Z_2 = x) \\ &= \sum_{y \in \mathbb{R}} \Pr(Z_1 = y) \Pr(Z_2 = x - y | Z_1 = y). \end{aligned} \tag{24}$$

Since the above convolution requires the calculation of conditional probability, it is computationally intractable. A common solution to it is to assume Z_1 and Z_2 to be mutually independent, which is a suitable assumption for decentralized utilities. This approach can be viewed as an approximation of the exact factorization method (Eq (24)). The approximation can be expressed as the following:

$$\begin{aligned} f(x) &= \sum_{y \in \mathbb{R}} \Pr(Z_1 = y) \Pr(Z_2 = x - y | Z_1 = y) \\ &\approx \sum_{y \in \mathbb{R}} \Pr(Z_1 = y) \Pr(Z_2 = x - y) \\ &= (f_1 * f_2)(x), \end{aligned} \tag{25}$$

where the star symbol (*) denotes the convolution operation. Eq (25) is tractable and allows C51 to model the PMF of each agent by a categorical distribution with n atoms (i.e., categories). The time complexity of deriving f for K agents is:

$$n \cdot n + n^2 \cdot n + \dots + n^{K-1} \cdot n = \sum_{k=1}^{K-1} n^{k+1} = O(n^K). \tag{26}$$

However, the time complexity grows exponentially as the number of the agents K increases. This becomes the training bottleneck and remains computationally infeasible when K is scaled to a large number. The computation cost can be reduced by performing an additional $O(n)$ heuristic projection after each convolution operation (Lin et al., 2019). The time complexity is then simplified to the following:

$$(n \cdot n + O(n)) \cdot (K - 1) = O(Kn^2). \tag{27}$$

The additional heuristic projections enables the decomposition process to be computationally feasible. However, such projections may induce an increase in variance for Z (Rowland

et al., 2019). If the return distribution of each agent shares the same support, the above $O(Kn^2)$ complexity can be improved by Fast Fourier Transform (FFT) convolutions. The technique is based on the convolution theorem, and is able to further reduce the time complexity to:

$$(n \log n + n + n \log n + O(n)) \cdot (K - 1) = O(Kn \log n). \quad (28)$$

4.3.2 SHAPE FUNCTION FOR IQN

Distributional RL methods such as IQN (Dabney et al., 2018a) do not explicitly model a distribution. They only implicitly model a distribution, and hence, the previously mentioned decomposition methods for C51 are not directly applicable. A potential solution is to add constraints on IQN to make the modeled distribution explicit (Yang et al., 2019), and iteratively apply convolution operations. To maintain the flexibility of IQN and to further reduce the time complexity, instead of imposing additional constraints, we propose to use quantile mixture to approximately factorize the distribution. In the following theorem, we prove that such a method is a valid factorization.

Theorem 3 *Given a quantile mixture:*

$$F^{-1}(\omega) = \sum_{k=1}^K \beta_k F_k^{-1}(\omega),$$

with K components $[F_k^{-1}]_{k \in \mathbb{K}}$, non-negative model parameters $[\beta_k]_{k \in \mathbb{K}}$, and $\omega \in [0, 1]$. There exist random variables Z and $[Z_k]_{k \in \mathbb{K}}$ derived from the quantile functions F^{-1} and $[F_k^{-1}]_{k \in \mathbb{K}}$, respectively, with the following relationship:

$$Z = \sum_{k \in \mathbb{K}} \beta_k Z_k.$$

Proof The proof is outlined as follows. First, we provide the notational details of the random variables. Next, we define the explicit forms of Z and $[Z_k]_{k \in \mathbb{K}}$. Finally, we prove the relationship $Z = \sum_{k \in \mathbb{K}} \beta_k Z_k$.

The domain of Z and $[Z_k]_{k \in \mathbb{K}}$ is Ω , where Ω is the sample space of the probability space (Ω, \mathcal{F}, P) . The codomain of Z and $[Z_k]_{k \in \mathbb{K}}$ is \mathbb{R} , where \mathbb{R} is the sample space of the probability spaces $(\mathbb{R}, \mathcal{B}, \mu)$ and $[(\mathbb{R}, \mathcal{B}, \mu_k)]_{k \in \mathbb{K}}$. Based on the definition of random variables in the measure theory, (Ω, \mathcal{F}, P) , $(\mathbb{R}, \mathcal{B}, \mu)$, and $[(\mathbb{R}, \mathcal{B}, \mu_k)]_{k \in \mathbb{K}}$ are defined as follows:

$$\begin{cases} \Omega = [0, 1], \text{ is the sample space,} \\ \mathcal{F}, \text{ is the Borel } \sigma\text{-algebra over } \Omega, \\ P : \mathcal{F} \rightarrow [0, 1], \text{ is a probability measure on the measurable space } (\Omega, \mathcal{F}). \\ \\ \mathbb{R} = (-\infty, \infty), \text{ is the real line,} \\ \mathcal{B}, \text{ is the Borel } \sigma\text{-algebra over } \mathbb{R}, \\ \mu : \mathcal{B} \rightarrow [0, 1], \text{ is a probability measure on the measurable space } (\mathbb{R}, \mathcal{B}), \\ [\mu_k : \mathcal{B} \rightarrow [0, 1]]_{k \in \mathbb{K}}, \text{ are probability measures on the measurable space } (\mathbb{R}, \mathcal{B}). \end{cases}$$

The random variables Z and $[Z_k]_{k \in \mathbb{K}}$ are defined as functions based on the probability spaces:

$$\begin{cases} Z : \Omega \rightarrow \mathbb{R}, \text{ with its CDF defined as } F(x) = \mu(-\infty, x], \\ [Z_k : \Omega \rightarrow \mathbb{R}]_{k \in \mathbb{K}}, \text{ with their CDFs defined as } F_k(x) = \mu_k(-\infty, x], \forall k \in \mathbb{K}. \end{cases}$$

We can then construct the random variables Z and $[Z_k]_{k \in \mathbb{K}}$ based on the components of the quantile mixture:

$$\begin{cases} Z(\omega) = F^{-1}(\omega), \forall \omega \in \Omega, \\ Z_k(\omega) = F_k^{-1}(\omega), \forall \omega \in \Omega, \forall k \in \mathbb{K}. \end{cases}$$

With the explicit forms of Z and $[Z_k]_{k \in \mathbb{K}}$, their relationship can be proved as follows:

$$\begin{aligned} F^{-1}(\omega) &= \sum_{k=1}^K \beta_k F_k^{-1}(\omega) \\ \Rightarrow Z(\omega) &= \sum_{k \in \mathbb{K}} \beta_k Z_k(\omega) \\ \Rightarrow Z &= \sum_{k \in \mathbb{K}} \beta_k Z_k \end{aligned}$$

■

In the optimization process, assume that all components $[F_k^{-1}]_{k \in \mathbb{K}}$ in the quantile mixture sample n quantiles, the total time complexity of summing the components is $O(Kn)$, which grows linearly with respect to K and n . The implementation is much easier and efficient when compared with the other methods described in the previous section. This suggests that quantile mixture is unlikely to become the training bottleneck, even if the values of K and n are both scaled to large numbers. It is worth noting that these two shape functions (i.e., convolution and quantile mixture) make different assumptions for modeling the joint return distribution.

This approach, however, trades off the exact factorization of the shape of the distribution for linear time complexity. Therefore, such a technique should only be used when the errors of capturing the shapes of the factorized distributions are acceptable. This technique is especially suitable for value function factorization methods since capturing the shapes of the factorized distributions is an auxiliary task.

4.4 DFAC Variant of VDN, QMIX, and QPLEX

In order to validate the proposed DFAC framework, we next discuss the DFAC variants of three representative factorization methods: VDN, QMIX, and QPLEX. For simplicity, these DFAC variants assume the shape function to be additive and non-learnable, and uses the quantile mixture technique instead of convolution. The DFAC variant of VDN is named DDN, which is expressed as follows:

$$Z_{\text{jt}} = \sum_{k \in \mathbb{K}} Q_k + \sum_{k \in \mathbb{K}} (Z_k - Q_k). \tag{29}$$

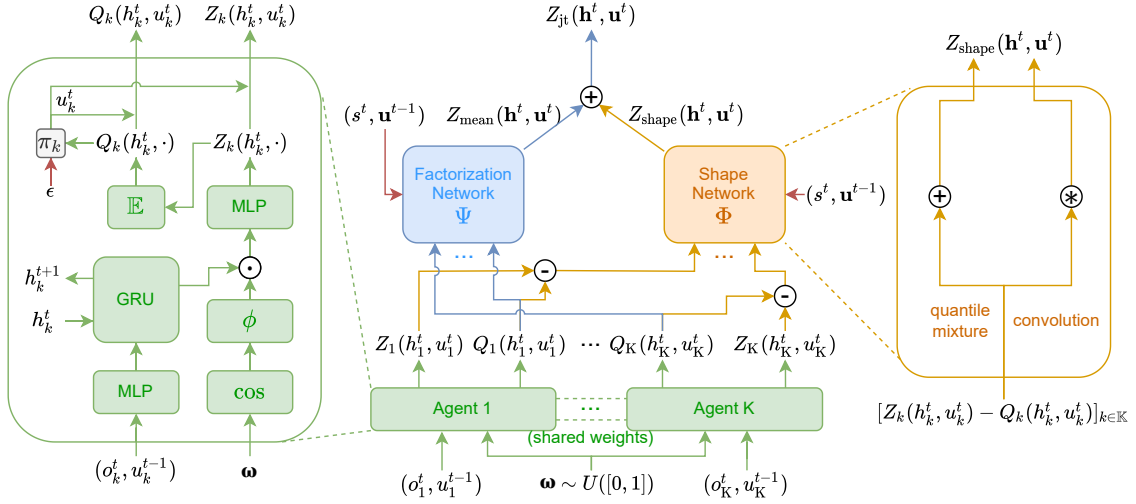


Figure 1: The DFAC framework consists of a factorization function Ψ and a shape function Φ for decomposing the deterministic part Z_{mean} (i.e., Q_{jt}) and the stochastic part Z_{shape} of the total return distribution Z_{jt} , as described in Theorem 2. If the shape function is a convolution, the network of the agents is implemented as C51. If the shape function is a quantile mixture, the network of the agents is implemented as IQN.

Eq (29) combines the factorization function of VDN (i.e., $\Psi = \sum_{k \in \mathbb{K}} Q_k$) and the additive shape function (i.e., $\Phi = \sum_{k \in \mathbb{K}} (Z_k - Q_k)$). In addition, the DFAC variant of QMIX is named DMIX, which is defined as the following:

$$Z_{jt} = M(Q_1, \dots, Q_K | s) + \sum_{k \in \mathbb{K}} (Z_k - Q_k). \quad (30)$$

Eq (30) combines the factorization function of QMIX (i.e., $\Psi = M(Q_1, \dots, Q_K | s)$), and the additive shape function (i.e., $\Phi = \sum_{k \in \mathbb{K}} (Z_k - Q_k)$). Lastly, the DFAC variant of QPLEX is named DPLEX, which is formulated as follows:

$$Q_{jt}(\mathbf{h}, \mathbf{u}) = \sum_{k=1}^K V_k(\mathbf{h}) + \sum_{k=1}^K \lambda_k(\mathbf{h}, \mathbf{u}) A_k(\mathbf{h}, u_k) + \sum_{k \in \mathbb{K}} (Z_k - Q_k). \quad (31)$$

Eq (31) is equivalent to summing the factorization function of QPLEX (i.e., $\Psi = \sum_{k=1}^K V_k(\mathbf{h}) + \sum_{k=1}^K \lambda_k(\mathbf{h}, \mathbf{u}) A_k(\mathbf{h}, u_k)$), and the additive shape function (i.e., $\Phi = \sum_{k \in \mathbb{K}} (Z_k - Q_k)$). These DFAC variants can be illustrated in the form of Fig. 1 with different factorization functions.

5. A Stochastic Matrix Game

In the previous expected value function factorization methods (e.g., VDN, QMIX, QPLEX, etc.), the factorization is achieved by modeling Q_{jt} and $[Q_k]_{k \in \mathbb{K}}$ as deterministic variables, overlooking the information of higher moments in the full return distributions Z_{jt} and

$[Z_k]_{k \in \mathbb{K}}$. In order to demonstrate DFAC’s ability of factorization, we begin with a toy example modified from Wang et al. (2021b) to show that DFAC is able to approximate the true return distributions, and factorize the mean and variance of the approximated total return Z_{jt} into utilities $[Z_k]_{k \in \mathbb{K}}$.

5.1 Experimental Setup

In this section, we describe the environment and methods, the architecture of the policy network, the architecture of the factorization network, and finally the training and execution details.

Environment and Methods. Table 1 illustrates a fully cooperative matrix game with stochastic rewards for two agents. The game only contains a single state, and each episode contains only a single step. In each episode, both agents choose an action and receive a global reward sampled from a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with mean μ and standard deviation σ . In this matrix game, we perform experiments on nine value function factorization methods: VDN, QMIX, and QPLEX are the baselines; DDN, DMIX, and DPLEX are the DFAC variants using IQN and quantile mixture; DDN-C51, DMIX-C51, and DPLEX-C51 correspond to the DFAC variants that use C51 and convolution. These methods are tuned independently and may have different hyperparameters.

Architecture of the Policy Network. Each agent’s policy network is implemented as an artificial neural network (ANN). QPLEX is implemented as a single-layered ANN with its hidden layer comprised of 32 units. DPLEX and DPLEX-C51 are implemented as two-layered ANNs comprised of 64 units and 32 units, with a ReLU nonlinearity at the end of the first layer. The other methods follow a similar architecture used by DPLEX and DPLEX-C51, where the second layer is changed to 512 units.

For DDN and DMIX, we optimize the IQNs of the agents with $N = N' = 32$ quantile samples, where each of them is first encoded into a 64-dimensional intermediate embedding and then projected to a 512-dimensional quantile embedding by a single hidden layer. As for DPLEX, we optimize the IQNs with $N = N' = 512$ quantile samples, where each of them is first encoded into a 256-dimensional intermediate embedding and then projected to a 32-dimensional quantile embedding by a single hidden layer.

For DDN-C51 and DMIX-C51, we optimize the C51 networks of the agents with 51 atoms uniformly distributed in $[-20, 20]$.

Architecture of the Factorization Network. For DMIX and its variants, we use a single-layered mixing network with eight units. As for QPLEX and its variants, we use a single-layered monotonic network with 16 units, and an attention network λ_k (i.e., in Eq (7)) with ten attention heads. Each of the attention heads utilizes a three-layered network with 64 units in the intermediate embeddings for its key, value, and query extractors.

Training and Execution. During training, each agent performs independent ϵ -greedy action selection, with full exploration (i.e., $\epsilon = 1$). The replay buffer contains experiences of the latest 2,000 episodes, from which we uniformly sample a batch of 2,048 samples when training QPLEX and its variants, and a batch of 512 samples for training the other six methods. The target network is updated every 100 episodes. The optimizer is set to Adam, in which its learning rate is set to 1×10^{-3} for QPLEX and QPLEX-C51, and 1×10^{-4} for the rest of the methods. We train each of the methods for 20,000 episodes. All of

Table 1: The payoff matrix of the stochastic matrix game. Each agent performs an action from $\{A, B, C\}$, with a subscript denoting the agent’s index. The global rewards are sampled from the normal distributions in the payoff matrix below.

		Agent 2		
		A_2	B_2	C_2
Agent 1	A_1	$\mathcal{N}(8, 8)$	$\mathcal{N}(-12, 6)$	$\mathcal{N}(-12, 4)$
	B_1	$\mathcal{N}(-12, 6)$	$\mathcal{N}(6, 4)$	$\mathcal{N}(0, 2)$
	C_1	$\mathcal{N}(-12, 4)$	$\mathcal{N}(0, 2)$	$\mathcal{N}(6, 0)$

Table 2: The value approximation errors and the returns of different methods.

Method	Q-dist	W-dist	Return
VDN	6.91	7.01	6
DDN	6.50	6.50	6
DDN-C51	6.01	6.17	6
QMIX	4.86	4.92	6
DMIX	3.76	4.58	6
DMIX-C51	4.96	5.29	6
QPLEX	0.10	1.47	8
DPLEX	0.06	0.24	8
DPLEX-C51	0.33	0.47	6

the agent networks share the same parameters. The agents are differentiated by a one-hot encoded agent index (i.e., $[1\ 0]^T$ for agent 1 and $[0\ 1]^T$ for agent 2) concatenated to their observations.

5.2 Experimental Results

Each column of Table 2 presents a key metric of the learned policy after convergence. *Q-dist* is defined as $\mathbb{E}[|Q_{jt} - Q_{jt}^*|]$, which is the averaged absolute distance between the approximated Q-value Q_{jt} and the true Q-value Q_{jt}^* under the optimal policy across all joint actions. *W-dist* represents $\mathbb{E}[W_1(Z_{jt}, Z_{jt}^*)]$, which is the averaged 1-Wasserstein distance between the approximated return Z_{jt} and the true return Z_{jt}^* under the optimal policy. *Return* reports the expectation of the returns following the learned policy. Each row of Table 2 corresponds to one of the methods mentioned in Section 5.1.

Based on the results, we observed that the methods with more expressiveness tend to have lower values of *Q-dist*. More specifically, QPLEX and its variants are able to model the expected return more precisely when compared to VDN, QMIX, and their variants. Moreover, the DFAC variants tend to have lower values of *W-dist*, reflecting DFAC’s capability of modeling the shapes of return distributions better than the baselines. Furthermore, QPLEX and DPLEX can learn the optimal return for this game (i.e., *Return* = 8), while VDN, QMIX, and their variants are unable to, which aligns with the results presented in Wang et al. (2021b). Unfortunately, DPLEX-C51 fails to learn the optimal return, even though it has the same theoretical expressiveness as QPLEX. We believe that this result is potentially due to the combination of C51’s shortcomings described in Section 3.5 and the convolutional operations used in the shape function for C51 mentioned in Section 4.3.1.

To further illustrate DFAC’s capability of factorization, we visualize the learned factorization of the joint values for QPLEX, DPLEX, and DPLEX-C51, respectively, in Figs. 2 and 3. These results demonstrate that QPLEX can only model the shapes of deterministic distributions, while its DFAC variants can model the shapes of both stochastic and deterministic distributions.

Since DPLEX-C51 does not perform satisfactorily in the matrix game (as indicated in Table 2), we only validate the IQN variants (DDN, DMIX, DPLEX) on the SMAC benchmark in the next section.

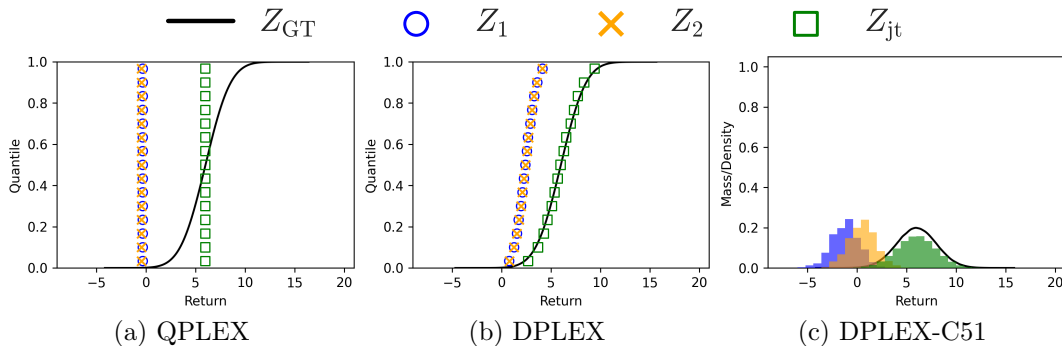


Figure 2: The learned factorization of the joint values when the joint action $\langle B_1, B_2 \rangle$ is selected. The ground truth (Z_{GT}) of the joint value is the stochastic return $\mathcal{N}(6, 4)$.




Figure 3: The learned factorization of the joint values when the joint action $\langle C_1, C_2 \rangle$ is selected. The ground truth (Z_{GT}) of the joint value is the deterministic return $\mathcal{N}(6, 0)$.

6. Experiment Results on SMAC

In this section, we present the experimental results and discuss their implications. We start with a brief introduction to our experimental setup in Section 6.1. Then, we demonstrate that modeling a full distribution is beneficial to the performance of independent learners in Section 6.2. Finally, we compare the performances of the CTDE baseline methods and their DFAC variants in Section 6.3 and Section 6.4.

6.1 Experimental Setup

In this section, we describe the environmental setups, the hyperparameter tuning process, as well as the baselines in details.

Environment. We verify the DFAC framework in the SMAC benchmark environments (Samvelyan et al., 2019) built on the popular real-time strategy game StarCraft II. Instead of playing the full game, SMAC is developed for evaluating the effectiveness of MARL micro-management algorithms. Each environment in SMAC contains two teams. One team is controlled by a decentralized MARL algorithm, with the policies of the agents conditioned on their local observation histories. The other team consists of enemy units controlled by the built-in game artificial intelligence based on carefully handcrafted heuristics, which is set to its highest difficulty equal to seven. The overall objective is to maximize the win

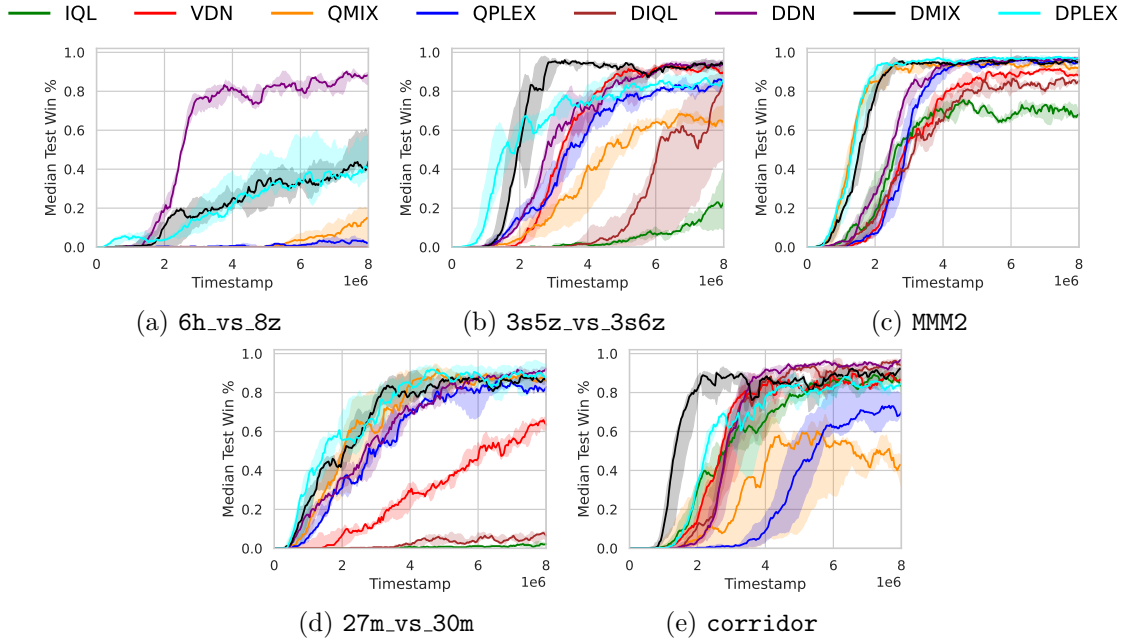


Figure 4: The win rate curves evaluated on the five *Super Hard* maps in SMAC.

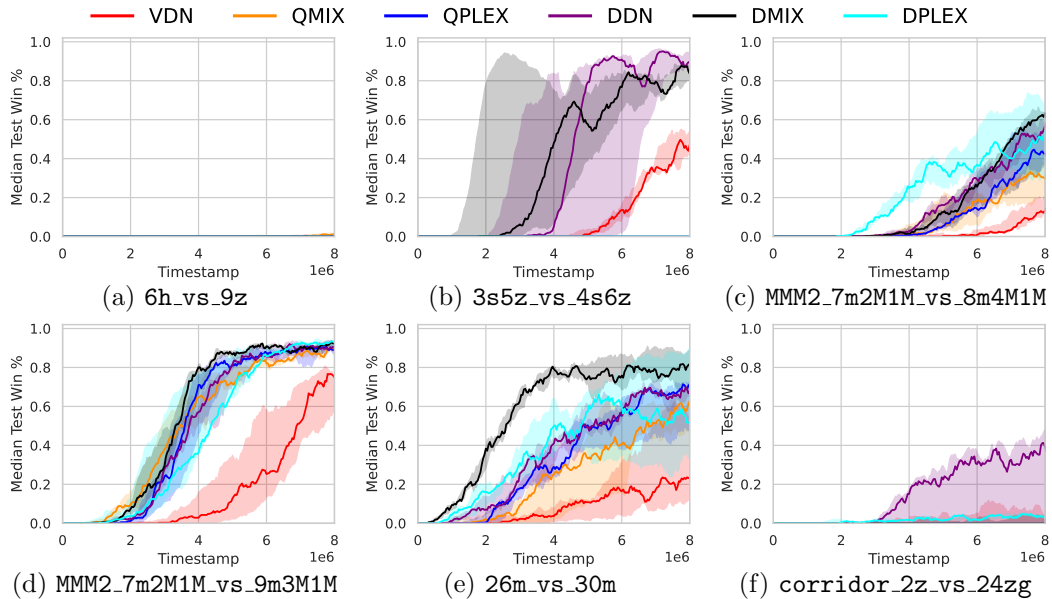


Figure 5: The win rate curves evaluated on the six *Ultra Hard* maps.

rate for each battle scenario, where the rewards employed in our experiments follow the default settings of SMAC. The default settings use *shaped rewards* based on the damage dealt, enemy units killed, and whether the RL agents win the battle. If there is no healing

unit in the enemy team, the maximum return of an episode (i.e., the score) is 20; otherwise, it may exceed 20, since enemies may receive more damages after healing or being healed.

The environments in SMAC are categorized into three different levels of difficulties: *Easy*, *Hard*, and *Super Hard* scenarios (Samvelyan et al., 2019). In this paper, we focus on all *Super Hard* scenarios including (a) `6h_vs_8z`, (b) `3s5z_vs_3s6z`, (c) `MMM2`, (d) `27m_vs_30m`, and (e) `corridor`, since these scenarios have not been properly addressed in the previous literature without the use of additional assumptions such as intrinsic reward signals (Du et al., 2019), explicit communication channels (Zhang et al., 2019; Wang et al., 2020b), common knowledge shared among the agents (de Witt et al., 2019; Wang et al., 2020a), and so on. Four of these scenarios have their maximum scores higher than 20. In `3s5z_vs_3s6z`, the enemy *Stalkers* have the ability to regenerate shields; in `MMM2`, the enemy *Medivacs* can heal other units; in `6h_vs_8z` and `corridor`, the enemy *Zerglings* slowly regenerate their own health. The StarCraft version we used is 4.10.

Hyperparameters. For all of our experimental results, the training length is set to 8M timesteps, where the agents are evaluated every 40k timesteps with 32 independent runs. The curves presented in this section are generated based on five different random seeds. The solid lines represent the median win rate, while the shaded areas correspond to the 25th to 75th percentiles. For a better visualization, the presented curves are smoothed by a moving average filter with its window size set to 11.

We tuned the hyperparameters of both the baselines and their distributional variants by selecting their hidden layer sizes from $\{32, 64, 128, 256, 512\}$ and choose the best ones. The quantile samples of DIQL (i.e., a distributional variant of IQL) and DDN are simply set to $N = N' = 1$, since they do not require the calculation of the expected value during the optimization process. As for DMIX and DPLEX, the numbers of quantile samples are set to $N = N' = 8$ as in Dabney et al. (2018a). The optimizers follow those used in DQN and IQN. All of the other hyperparameters follow those used in SMAC. Table 6 lists the hyperparameters adopted for the baselines and their distributional variants.

Baselines. We select IQL, VDN, QMIX, and QPLEX as our baseline methods, and compare them with their distributional variants in our experiments. The configurations are optimized so as to provide the best performance for each of the methods considered. Since we tuned the hyperparameters of the baselines, their performances are better than those reported in (Samvelyan et al., 2019). The hyperparameter searching process is detailed in the supplementary material.

6.2 Independent Learners

In order to validate our assumption that distributional RL is beneficial to the MARL domain, we first employ the simplest training algorithm, IQL, and extend it to its distributional variant, called DIQL. DIQL is simply a modified IQL that uses IQN as its underlying RL algorithm without any additional modification or enhancements (Maignon et al., 2007; Lyu and Amato, 2020).

From Figs. 4(a)-4(e) and Tables 3 and 4, it is observed that DIQL is superior to IQL even without utilizing any value function factorization methods. This validates that distributional RL has beneficial influences on MARL, when it is compared to RL approaches based only on expected values.

6.3 Value Function Factorization Methods

In order to inspect the effectiveness and impacts of DFAC on learning curves, win rates, and scores, we next summarize the results of the baselines as well as their DFAC variants on the *Super Hard* scenarios in Fig. 4(a)-(e) and Tables 3 and 4.

Fig. 4(a)-(e) plot the learning curves of the baselines and their DFAC variants, with the final win rates presented in Table 3, and their final scores reported in Table 4. The win rates indicate how often do the player’s team wins, while the scores represent how well do the player’s team performs. Despite the fact that SMAC’s objective is to maximize the win rate, the true optimization goal of MARL algorithms is the averaged score. In fact, these two metrics are not always positively correlated (e.g., VDN and QMIX in *6h_vs_8z* and *3s5z_vs_3s6z*, and QMIX and DMIX in *3s5z_vs_3s6z*).

It can be observed that the learning curves of DDN, DMIX, and DPLEX grow faster and achieve higher final win rates than their corresponding baselines. In the most difficult map: *6h_vs_8z*, most of the methods fail to learn an effective policy except for the DFAC variants. The evaluation results also show that the DFAC variants are capable of performing consistently well across all *Super Hard* maps with high win rates. In addition to the win rates, Table 4 further presents the final averaged scores achieved by each method, and provides deeper insights into the advantages of the DFAC framework by quantifying the performances of the learned policies of different methods.

The improvements in win rates and scores are due to the benefits offered by distributional RL (Lyle et al., 2019), which enables the distributional variants to work more effectively in MARL environments. Moreover, the evaluation results reveal that DDN performs especially well in most environments despite its simplicity.

6.4 Evaluation Results on the Ultra Hard Maps of SMAC

Based on the five *Super Hard* maps, we further designed six *Ultra Hard* maps for evaluating the capability of DFAC. The *Ultra Hard* maps are guaranteed to be harder than the *Super Hard* maps, since they contain either increased number of units in the enemy team or decreased number of controllable units in the player’s team. The detailed comparison between the two sets of maps is presented in Table 5.

The evaluation results on the *Ultra Hard* maps are reported in Fig. 5 and Tables 3 and 4 with five independent runs. It is observed that DDN achieves outstanding performance in most of the maps despite the simplicity of its factorization function. For reproducibility, all *Ultra Hard* maps can be found in our GitHub repository (<https://github.com/j3soon/dfac-extended>), along with the gameplay recording videos.

6.5 Computational Infrastructure

In our experiments, NVIDIA DGX-1 clusters were used as the fundamental computing infrastructure, where each DGX-1 contains eight V100 GPUs, two 20-core CPUs, and 512 GB memory. Our experiments were performed on docker instances, where each instance is allocated with 50 GB DDR4 memory, an eight-core Xeon E5-2698 CPU, 1 TB SSD storage, and either a V100-16G or a V100-32G GPU. On an average, the experiment for each map

Table 3: The median win rate percentage (%) of five independent test runs.

	Map	IQL	VDN	QMIX	QPLEX	DIQL	DDN	DMIX	DPLEX
<i>Super Hard</i>	6h_vs_8z	0.00	0.00	12.78	0.00	0.00	83.92	49.43	43.75
	3s5z_vs_3s6z	29.83	89.20	67.22	84.38	62.22	94.03	91.08	90.62
	MMM2	68.92	89.20	92.44	96.88	85.23	97.22	95.11	96.88
	27m_vs_30m	2.27	63.12	84.77	78.12	6.02	91.48	85.45	90.62
	corridor	84.87	85.34	37.61	75.00	91.62	95.40	90.45	81.25
<i>Ultra Hard</i>	6h_vs_9z	-	0.00	1.14	0.00	-	0.28	0.00	0.00
	3s5z_vs_4s6z	-	47.16	0.00	0.00	-	89.77	83.52	0.00
	MMM2_7m2M1M_vs_8m4M1M	-	13.35	29.55	46.88	-	56.82	63.35	50.00
	MMM2_7m2M1M_vs_9m3M1M	-	75.00	88.64	90.62	-	90.34	92.33	90.62
	26m_vs_30m	-	23.01	62.78	78.12	-	67.90	81.82	59.38
	corridor_2z_vs_24zg	-	0.00	0.00	0.00	-	41.19	0.00	3.12

Table 4: The averaged scores of five independent test runs.

	Map	IQL	VDN	QMIX	QPLEX	DIQL	DDN	DMIX	DPLEX
<i>Super Hard</i>	6h_vs_8z	13.78	15.41	14.37	15.95	14.94	19.40	17.14	17.88
	3s5z_vs_3s6z	16.54	19.75	20.16	20.42	17.52	20.94	19.70	20.27
	MMM2	17.50	19.36	19.42	19.60	19.21	20.90	19.87	19.93
	27m_vs_30m	14.01	18.45	19.41	19.33	14.45	19.71	19.43	19.62
	corridor	19.42	19.47	15.07	18.73	19.68	20.00	19.66	19.08
<i>Ultra Hard</i>	6h_vs_9z	-	13.57	12.37	13.86	-	16.00	13.73	14.84
	3s5z_vs_4s6z	-	17.16	13.09	13.60	-	19.65	18.61	14.99
	MMM2_7m2M1M_vs_8m4M1M	-	13.13	14.40	15.52	-	16.50	16.24	15.89
	MMM2_7m2M1M_vs_9m3M1M	-	17.30	19.01	19.06	-	19.45	19.33	19.40
	26m_vs_30m	-	16.69	18.23	18.66	-	18.49	19.17	18.49
	corridor_2z_vs_24zg	-	7.78	4.80	6.44	-	11.10	7.41	10.71

takes around forty hours to train for eight million timesteps, depending on the type of the SMAC map and the chosen algorithm.

7. Discussions and Outlook

In this section, we provide discussions on the SMAC benchmark and propose some future directions to extend DFAC.

The Performance Metric. In the original SMAC paper, the authors proposed to use the median test win rates as the main performance metric. However, their optimization goal is actually the discounted scores instead of the win rates. Therefore, the averaged scores of multiple test runs are better metrics for comparing the performance of different methods.

Enemy Sight Range. By inspecting the learned policies of DDN, we observed that the agents can exploit enemies’ limited sight ranges by luring out and eliminating a few enemies while keeping distances from the remainings in `corridor_2z_vs_24zg`. To make

Table 5: A comparison between the *Super Hard* maps and the *Ultra Hard* maps.

	Map	Player’s Team	Enemy’s Team
<i>Super Hard</i>	6h_vs_8z	6 Hydralisks	8 Zealots
	3s5z_vs_3s6z	3 Stalkers & 5 Zealots	3 Stalkers & 6 Zealots
	MMM2	7 Marines, 2 Marauders & 1 Medivac	8 Marines, 3 Marauders & 1 Medivac
	27m_vs_30m	27 Marines	30 Marines
	corridor	6 Zealots	24 Zerglings
<i>Ultra Hard</i>	6h_vs_9z	6 Hydralisks	9 Zealots
	3s5z_vs_4s6z	3 Stalkers & 5 Zealots	4 Stalkers & 6 Zealots
	MMM2_7m2M1M_vs_8m4M1M	7 Marines, 2 Marauders & 1 Medivac	8 Marines, 4 Marauders & 1 Medivac
	MMM2_7m2M1M_vs_9m3M1M	7 Marines, 2 Marauders & 1 Medivac	9 Marines, 3 Marauders & 1 Medivac
	26m_vs_30m	26 Marines	30 Marines
	corridor_2z_vs_24zg	2 Zealots	24 Zerglings

the map more challenging, the enemies can be configured to have infinite sight ranges when designing the map.

The Reward Function. The reward function is defined as the accumulated damages dealt to the enemies. However, the enemies may regenerate their health and shields. This may lead to a scenario that the agents intentionally stop attacking the enemies with low health and attack them later after their health is regenerated, allowing the agents to collect more rewards. Such a behavior can be prevented by providing no reward signal for regenerated health.

Expressiveness and Performance. According to the results presented in Section 6, the *Super Hard* maps can be solved under the assumption of *Additivity*, since DDN is able to achieve high win rates on all of these maps. These results indicate that the other methods with additional expressiveness have no theoretical advantage over DDN on these *Super Hard* maps, and may perform worse due to the instability caused by their factorization functions. This further verifies the fact that the expressiveness is not necessarily correlated with the performance in SMAC, which is often misinterpreted in the current literature.

Network Size and Training Steps. In the original SMAC paper, the authors did not tune the size of the policy networks and only trained them for two million timesteps. According to our experiments, we found that using larger policy networks and training them for more timesteps does improve the performance by a noticeable margin.

Future Extensions. Based on the DFAC framework, many SARL techniques that require the shapes of the return distributions may be incorporated into the MARL domain, such as exploration (Nikolov et al., 2019; Zhang and Yao, 2019; Mavrin et al., 2019) and risk-aware policies (Xia, 2020). Aside from the simple shape function used in this work, future endeavors may include more complex shape functions with learnable weights. It would also be interesting to incorporate other distributional RL algorithms into the DFAC framework aside from C51 and IQN, such as quantile regression DQN (QR-DQN) (Dabney et al., 2018b), expectile distributional RL (EDRL) (Rowland et al., 2019), fully parameterized quantile function (FQF) (Yang et al., 2019), and maximum mean discrepancy RL (MMDRL) (Nguyen et al., 2021).

Table 6: A summary of the optimal hidden state sizes of the baseline methods and their distributional variants.

	Map	IQL	VDN	QMIX	QPLEX	DIQL	DDN	DMIX	DPLEX
<i>Super Hard</i>	6h_vs_8z	128	128	256	256	512	512	256	512
	3s5z_vs_3s6z	512	128	128	128	256	512	256	512
	MMM2	256	64	64	64	512	512	256	256
	27m_vs_30m	256	64	64	64	512	128	128	128
	corridor	256	128	256	64	512	128	64	512
<i>Ultra Hard</i>	6h_vs_9z	-	256	64	256	-	128	512	256
	3s5z_vs_4s6z	-	64	64	128	-	512	512	256
	MMM2_7m2M1M_vs_8m4M1M	-	128	128	128	-	256	128	256
	MMM2_7m2M1M_vs_9m3M1M	-	64	64	128	-	256	128	128
	26m_vs_30m	-	32	32	64	-	128	64	128
	corridor_2z_vs_24zg	-	128	512	256	-	256	512	512

8. Conclusion

In this paper, we introduced a unified framework for cooperative MARL, called DFAC, for integrating distributional RL with value function factorization. DFAC is based on mean-shape decomposition to meet the Distributional IGM condition. To realize DFAC in a computationally friendly manner, its shape function is implemented as convolution for C51, and quantile mixture for IQN. DFAC’s ability to factorize a joint return distribution into individual utility distributions was demonstrated in a matrix game. In order to validate the effectiveness of DFAC, we presented experimental results performed on all *Super Hard* scenarios in SMAC for a number of MARL baseline methods as well as their DFAC variants. Moreover, we performed experiments on a number of self-designed *Ultra Hard* maps to further validate the effectiveness of DFAC. The results showed that in most of the scenarios, DDN, DMIX, and DPLEX outperform VDN, QMIX, and QPLEX, respectively. DFAC can be extended to other value function factorization methods and offers an interesting research direction for future endeavors.

Acknowledgments

The authors gratefully acknowledge the support from the National Science and Technology Council (NSTC) in Taiwan under grant number MOST 111-2223-E-007-004-MY3. The authors would also like to express their appreciation for the donation of the GPUs from NVIDIA Corporation and NVIDIA AI Technology Center (NVAITC) used in this work. In addition, the authors extend their gratitude to the National Center for High-Performance Computing (NCHC) for providing the necessary computational and storage resources. Furthermore, the authors thank Kuan-Yu Chang for his helpful critiques of this research work.

References

- M. G. Bellemare, W. Dabney, and R. Munos. A distributional perspective on reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 449–458, Jul. 2017.
- M. G. Bellemare, N. Le Roux, P. S. Castro, and S. Moitra. Distributional reinforcement learning with linear function approximation. In *Proc. Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, pages 2203–2211, Apr. 2019.
- M. G. Bellemare, W. Dabney, and M. Rowland. *Distributional Reinforcement Learning*. MIT Press, 2022. <http://www.distributional-rl.org>.
- F. L. Da Silva, A. H. R. Costa, and P. Stone. Distributional reinforcement learning applied to robot soccer simulation. In *Workshop on Adaptive Learning Agents (ALA) at AAMAS*, Jun. 2019.
- W. Dabney, G. Ostrovski, D. Silver, and R. Munos. Implicit quantile networks for distributional reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 1096–1105, Jul. 2018a.
- W. Dabney, M. Rowland, M. G. Bellemare, and R. Munos. Distributional reinforcement learning with quantile regression. In *Proc. AAAI Conf. on Artificial Intelligence (AAAI)*, pages 2892–2901, Feb. 2018b.
- C. S. de Witt et al. Multi-agent common knowledge reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 9924–9935, Dec. 2019.
- Y. Du et al. LIIR: Learning individual intrinsic reward in multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4405–4416, Dec. 2019.
- J. N. Foerster, G. Farquhar, T. Afouras, N. Nardelli, and S. Whiteson. Counterfactual multi-agent policy gradients. In *Proc. AAAI Conf. on Artificial Intelligence (AAAI)*, Feb. 2018.
- C. Guestrin, D. Koller, and R. Parr. Multiagent planning with factored MDPs. In *Advances in Neural Information Processing Systems (NIPS)*, 2001.
- J. Hu, S. A. Harding, H. Wu, and S. W. Liao. QR-MIX: Distributional value function factorisation for cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2009.04197*, 2020.
- S. Iqbal and F. Sha. Actor-attention-critic for multi-agent reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 2961–2970, Jul. 2019.
- J. Karvanen. Estimation of quantile mixtures via L-moments and trimmed L-moments. *Computational Statistics & Data Analysis*, pages 947–959, Nov. 2006.
- Z. Lin et al. Distributional reward decomposition for reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6212–6221, Dec. 2019.

- R. Lowe et al. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6379–6390, Dec. 2017.
- C. Lyle, M. G. Bellemare, and P. S. Castro. A comparative analysis of expected and distributional reinforcement learning. In *Proc. AAAI Conf. on Artificial Intelligence (AAAI)*, pages 4504–4511, Feb. 2019.
- X. Lyu and C. Amato. Likelihood quantile networks for coordinating multi-agent reinforcement learning. In *Proc. Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 798–806, May 2020.
- L. Matignon, G. Laurent, and N. Fort-Piat. Hysteretic Q-Learning: an algorithm for decentralized reinforcement learning in cooperative multi-agent teams. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 64–69, Dec. 2007.
- B. Mavrin, H. Yao, L. Kong, K. Wu, and Y. Yu. Distributional reinforcement learning for efficient exploration. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 4424–4434, Jul. 2019.
- V. Mnih et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, Feb. 2015.
- T. T. Nguyen, S. Gupta, and S. Venkatesh. Distributional reinforcement learning via moment matching. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Feb. 2021.
- N. Nikolov, J. Kirschner, F. Berkenkamp, and A. Krause. Information-directed exploration for deep reinforcement learning. In *Proc. Int. Conf. on Learning Representations (ICLR)*, May 2019.
- F. A. Oliehoek and C. Amato. *A Concise Introduction to Decentralized POMDPs*. Springer, 2016.
- S. Omidshafiei, J. Pazis, C. Amato, J. P. How, and J. Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 2681–2690, Aug. 2017.
- B. Peng et al. FACMAC: Factored multi-agent centralised policy gradients. In *Advances in Neural Information Processing Systems (NIPS)*, pages 12208–12221, Dec. 2021.
- W. Qiu et al. RMIX: Learning risk-sensitive policies for cooperative reinforcement learning agents. *Advances in Neural Information Processing Systems (NIPS)*, pages 23049–23062, Dec. 2021.
- N. Rahaman et al. On the spectral bias of neural networks. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 5301–5310, Jul. 2019.
- T. Rashid, G. Farquhar, B. Peng, and S. Whiteson. Weighted QMIX: Expanding monotonic value function factorisation. *Advances in Neural Information Processing Systems (NIPS)*, pages 10199–10210, Dec. 2020.

- T. Rashid et al. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 4295–4304, Jul. 2018.
- M. Rowland et al. Statistics and samples in distributional reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 5528–5536, Jul. 2019.
- M. Rowland et al. Temporal difference and return optimism in cooperative multi-agent reinforcement learning. In *Workshop on Adaptive Learning Agents (ALA) at AAMAS*, May 2021.
- M. Samvelyan et al. The StarCraft multi-agent challenge. In *Proc. Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 2186–2188, May 2019.
- T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 1312–1320, Jul. 2015.
- K. Son, D. Kim, W. J. Kang, D. E. Hostallero, and Y. Yi. QTRAN: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 5887–5896, Jul. 2019.
- W. F. Sun, C. K. Lee, and C. Y. Lee. DFAC framework: Factorizing the value function via quantile mixture for multi-agent distributional Q-Learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 9945–9954, Jul. 2021.
- P. Sunehag et al. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proc. Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 2085–2087, May 2018.
- M. Tan. Multi-agent reinforcement learning: Independent versus cooperative agents. In *Proc. Int. Conf. on Machine Learning (ICML)*, page 330–337, Jun. 1993.
- J. Wang, Z. Ren, B. Han, J. Ye, and C. Zhang. Towards understanding cooperative multi-agent Q-Learning with value factorization. *Advances in Neural Information Processing Systems (NIPS)*, Dec. 2021a.
- J. Wang, Z. Ren, T. Liu, Y. Yu, and C. Zhang. QPLEX: Duplex dueling multi-agent Q-Learning. In *Proc. Int. Conf. on Learning Representations (ICLR)*, May 2021b.
- T. Wang, H. Dong, V. Lesser, and C. Zhang. Multi-agent reinforcement learning with emergent roles. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 9876–9886, Jul. 2020a.
- T. Wang, J. Wang, C. Zheng, and C. Zhang. Learning nearly decomposable value functions via communication minimization. In *Proc. Int. Conf. on Learning Representations (ICLR)*, Apr. 2020b.
- W. Wang et al. Action semantics network: Considering the effects of actions in multiagent systems. In *Proc. Int. Conf. on Learning Representations (ICLR)*, Apr. 2020c.

- L. Xia. Risk-sensitive Markov decision processes with combined metrics of mean and variance. *Production and Operations Management*, 29(12):2808–2827, 2020.
- D. Yang et al. Fully parameterized quantile function for distributional reinforcement learning. In *Proc. Conf. Advances in Neural Information Processing Systems (NeurIPS)*, pages 6190–6199, Dec. 2019.
- Y. Yang, J. Hao, B. Liao, K. Shao, G. Chen, W. Liu, and H. Tang. Qatten: A general framework for cooperative multiagent reinforcement learning. *arXiv preprint arXiv:2002.03939*, 2020.
- S. Zhang and H. Yao. QUOTA: The quantile option architecture for reinforcement learning. In *Proc. AAAI Conf. on Artificial Intelligence (AAAI)*, pages 5797–5804, Feb. 2019.
- S. Q. Zhang, Q. Zhang, and J. Lin. Efficient communication in multi-agent reinforcement learning via variance based control. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3230–3239, Dec. 2019.