

# Tackling Uncertainties in Multi-Agent Reinforcement Learning through Integration of Agent Termination Dynamics

Somnath Hazra  
IIT Kharagpur  
Kharagpur, India  
somnathazra@kgpian.iitkgp.ac.in

Pallab Dasgupta  
Synopsis  
Santa Clara, USA  
pallabd@synopsys.com

Soumyajit Dey  
IIT Kharagpur  
Kharagpur, India  
soumya@cse.iitkgp.ac.in

## ABSTRACT

Multi-Agent Reinforcement Learning (MARL) has gained significant traction for solving complex real-world tasks, but the inherent stochasticity and uncertainty in these environments pose substantial challenges to efficient and robust policy learning. While Distributional Reinforcement Learning has been successfully applied in single-agent settings to address risk and uncertainty, its application in MARL is substantially limited. In this work, we propose a novel approach that integrates distributional learning with a safety-focused loss function to improve convergence in cooperative MARL tasks. Specifically, we introduce a Barrier Function based loss that leverages safety metrics, identified from inherent faults in the system, into the policy learning process. This additional loss term helps mitigate risks and encourages safer exploration during the early stages of training. We evaluate our method in the StarCraft II micromanagement benchmark, where our approach demonstrates improved convergence and outperforms state-of-the-art baselines in terms of both safety and task completion. Our results suggest that incorporating safety considerations can significantly enhance learning performance in complex, multi-agent environments.

## CCS CONCEPTS

• **Theory of computation** → **Multi-agent reinforcement learning**; *Multi-agent learning*.

## KEYWORDS

Multi-Agent Reinforcement Learning, Distributional RL, Barrier Function

## ACM Reference Format:

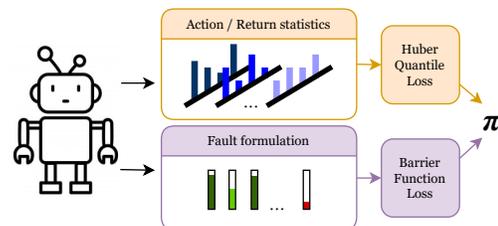
Somnath Hazra, Pallab Dasgupta, and Soumyajit Dey. 2025. Tackling Uncertainties in Multi-Agent Reinforcement Learning through Integration of Agent Termination Dynamics. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

## 1 INTRODUCTION

Multi-Agent Systems (MAS) remain a highly challenging domain, owing to the inherent stochasticity in the environment, and concurrently learning agents; resulting in uncertain reward outcomes. Therefore the contemporary Multi-Agent Reinforcement Learning (MARL) literature has shifted from point estimates derived from

expected global rewards [24, 36] to distributional analysis over returns [4, 8] which provide a more informed understanding of the variability in future returns. However, the predicted distributions are prone to errors at the initial phases of training because the policy has not yet explored and trained enough to predict accurate distributions. In the context of MARL, the prediction errors become even more critical due to the collective uncertainty from multiple agents, where each agent’s decisions can drastically impact both individual and collective outcomes. Our focus lies specifically on cooperative MARL, where a team of agents collaborates to achieve a shared objective in a stochastic environment [1].

Despite the effectiveness of the Centralized Training and Decentralized Execution (CTDE) paradigm [16] for cooperative MARL most CTDE-based algorithms maximize returns and overlook the uncertainty due to untrained parameters in the policy that captures the inherent stochasticity in MAS. Here agents must contend with not only uncertain rewards but also additional safety constraints resulting from the stochasticity. Our contribution is to identify the innate fault tolerant MARL formulation from the system such that the system does not go beyond a pre-defined safety boundary. For example, a team game cannot be won if players keep getting eliminated. These safety objectives hold importance during the early phases of training since it is hard to estimate accurate distributions by a distributional model owing to exploration strategies and partial observability of the agents.



**Figure 1: Integration of the safety constraints along with loss functions calculated with respect to returns, helps reduce the uncertainty, especially during the initial phase of training.**

As an instance of such a situation, let us consider a battle scenario between two agent groups. The reward can be based on intermediate strategies decided by the teams. But it is unsure whether their strategy will work, since any agent army does not know about the strategies planned by the other army. Moreover the agents may not be experts in carrying out their decided objective. However, one property that is overlooked, is that the battle cannot be won if there are not enough agents left to fight; i.e., there is an inherent fault tolerant formulation in the system. For example, in battle



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

simulations like StarCraft II, minimizing ally casualties can be as important as maximizing reward [26].

Motivated by the gap in current research, our work addresses the uncertainties of the system via integration of safety considerations from the inherent faults in the MAS. While previous efforts have successfully combined Distributional RL with MARL to improve risk sensitivity [23, 33], these approaches largely focus on algorithmic modifications rather than optimizing the learning process by using the available information regarding the safety-critical objectives. Our contributions are summarized as follows:

- We propose a novel loss function that incorporates safety constraints into the optimization process, derived from the safety boundary of the inherent faults in the system. By explicitly accounting for safety alongside reward maximization, we aim to accelerate learning, tackle environmental uncertainties, and improve long-term policy performance in risk-sensitive environments. We also use the Dueling Networks [39] in our architecture to eliminate ineffective actions.
- We also suggest an architectural improvement for the local policy network in the CTDE paradigm to utilize the return distribution for prioritizing the important components from the local observation.
- We validated our approach on the StarCraft II micromanagement [26] and MetaDrive[15] benchmarks, demonstrating its efficacy compared to existing baselines and state-of-the-art distributional MARL algorithms.

More specifically, our work contributes to bridging the gap between theoretical advancements in Distributional MARL and practical needs in real-world MARL environments, particularly those with critical safety concerns.

## 2 PRELIMINARIES

This section outlines some of the preliminary concepts.

### 2.1 Decentralized POMDP

The de facto standard for modeling cooperation in MARL is a Decentralized Partially Observable MDP or Dec-POMDP [21]. It is defined using the tuple  $\mathcal{M} = \langle \mathcal{N}, \mathcal{S}, \mathcal{U}, \Omega, P, O, r, \gamma \rangle$ , where  $\mathcal{N} \equiv \{a_1, a_2, \dots, a_n\}$  is a finite set of agents, and  $s \in \mathcal{S}$  represents the global state of the system. At each time-step, the individual agents choose an action  $u_i \in \mathcal{U}$  for  $i \in \mathcal{N}$  based on observation  $o_i \in \Omega$ , with probability  $\mathcal{O}(o_i | s, u_i)$ . The joint action,  $\mathbf{u} = [u_i]_{i=1}^n$ , changes the state of the system from  $s$  to  $s' \sim P(\cdot | s, \mathbf{u})$ ; yielding a joint reward  $r(s, \mathbf{u})$ .  $\gamma \in [0, 1)$  is the discount factor.

Value functions provide an estimate of the reward that can be obtained given a state, or when taking an action from a state. Accordingly, the state value function for the system is given by,  $V_{\pi}(s) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{u}_t) | s_0 = s, \pi]$ ; and the state-action value function is given by,  $Q_{\pi}(s, \mathbf{u}) = r(s, \mathbf{u}) + \gamma \mathbb{E}_{s'} V_{\pi}(s')$ . The formal objective of the co-operative system is to find a joint policy  $\pi \equiv \langle \pi_1, \dots, \pi_n \rangle$  that maximizes the expected return given by:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{(s_t, \mathbf{u}_t) \sim (P, \pi)} \left[ \sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{u}_t) \right] \quad (1)$$

### 2.2 Centralized Training and Decentralized Execution

Multi-agent algorithms generally follow the Centralized Training Decentralized Execution (CTDE) paradigm [16]. Agents learn collaboratively during training using feedback from the environment. This phase may involve information exchange with a central approximation network. This phase is termed as the Central Training phase. The Central Value function accumulates the values of the individual policies  $\pi_i$  in a forward propagation determined by the *Credit Assignment* strategy [7] of the algorithm. The strategies build on the *Individual Global Max (IGM)* principle which states that locally optimal actions maximize the global reward of the system given by the following equation [12].

$$\arg \max_{\mathbf{u} \in \mathcal{U}^n} Q_{\pi}(s, \mathbf{u}) = \left[ \arg \max_{u_i \in \mathcal{U}} Q_{\pi_i}(o_i, u_i) \right]_{i=1}^n \quad (2)$$

During the Decentralized Execution phase, agents act independently based on local observations.

### 2.3 Distributional RL

In single-agent RL, distributional RL aims to learn the complete distribution over the returns,  $Z(s, \mathbf{u})$ , instead of learning the expected return,  $Q(s, \mathbf{u})$ . The distribution over the returns can be approximated using a categorical distribution [4] or a quantile function [8]. We here use Implicit Quantile Network (IQN) [8], that models  $Z(s, \mathbf{u})$  using an inverse cumulative distribution function (CDF),  $F^{-1}(s, \mathbf{u} | w)$ ,  $w \in [0, 1]$ . IQN defines the distributional Bellman operator to update  $Z(s, \mathbf{u})$  using the Huber Quantile Regression loss, a distributional version of the temporal difference error. During execution, the action is chosen based on the largest expected return,  $\arg \max_{\mathbf{u}} \mathbb{E}[Z(s, \mathbf{u})]$ . The distributional analysis in MARL is based on the *Distributional IGM (DIGM)* principle, which states that, given individual state-action value distributions,  $[Z_{\pi_i}(o_i, u_i)]_{i=1}^n$ , the joint state-action value distribution,  $Z_{\pi}(s, \mathbf{u})$ , is maximized if the following equation holds [33].

$$\arg \max_{\mathbf{u} \in \mathcal{U}^n} Z_{\pi}(s, \mathbf{u}) = \left[ \arg \max_{u_i \in \mathcal{U}} Z_{\pi_i}(o_i, u_i) \right]_{i=1}^n \quad (3)$$

## 3 RELATED WORK

Efficiency and scalability have been central challenges in the MARL literature, particularly when estimating the global expected return (for using the reward,  $r$  as a reference during training) from local policies (in decentralized execution). Value factorization methods have been widely adopted to address these challenges, adhere to the IGM principle. Initial approaches like VDN [36] directly add local value functions; while QMIX [24] uses a monotonic mixing function to ensure compatibility with the global return. QTRAN [30] introduced linear constraints to factorize the global return without the monotonicity assumption. Transformer-based architectures, such as Qatten [40], have also been explored in MARL. QPLEX [38] leverages the dueling network architecture [39] to improve learning generalization across actions, which we have used in a distributional setting in this work. Despite these advancements, the stochasticity inherent in multi-agent systems complicates the use of expected returns, often hindering training efficiency.

In parallel, distributional RL has made significant strides in the single-agent domain, introducing algorithms such as C51 [4], Implicit Quantile Networks (IQN) [8], and others [9, 25]. More recently, there have been increasing interest in unifying distributional RL with MARL to improve risk-sensitivity and robustness. Notable works include restructuring the IGM principle to a distributional Q-learning standpoint [27, 33]; or exploring the methods for aggregating reward distributions for each action, considering the sources of risk [20, 32]. Works such as RMIX [23] integrate risk-aware metrics like Conditional Value at Risk (CVaR) into the QMIX framework, while Sun et. al. [33] introduced the Distributional IGM (DIGM) principle, extending value factorization methods like VDN and QMIX to the distributional setting. DIGM was further refined in subsequent work [35]. ResQ [28] builds upon DMIX by introducing residual Q-functions for more accurate return estimation, while other studies have explored risk-sensitive aggregation of reward distributions [20, 32], uncertainty-aware exploration strategies [19], and modelling uncertainties related to reward [13]. Our work builds on these existing distributional MARL approaches by incorporating agent-specific information to enhance policy learning; unlike previous studies that focus primarily on algorithmic adaptations of MARL to the distributional context. This enables more effective training in stochastic multi-agent environments.

Action shielding has been another widely explored venue in RL and MARL to prevent unsafe actions during training and execution. Alshiekh et al. [2] introduced a shielding mechanism to enforce safety constraints by overriding unsafe actions based on a safety specification. Building on this, Bharadwaj et al. [5] proposed minimum-cost shields for multi-agent systems to ensure safe coordination. Factored shields curated on linear temporal logic properties [10] is another venue that has been explored for MARL. For environments with non-linear dynamics, model predictive shielding [3] provides an efficient approach to maintaining safety. Zhang et al. proposed a multi-agent version [43], where only a subset of agents, determined by a greedy algorithm, need to use a backup policy. While these methods focus on action-level intervention, our approach incorporates safety directly into the policy learning process via the barrier function, complementing action shielding techniques by prioritizing long-term safety considerations during training.

## 4 METHODOLOGY

The dynamic nature of MARL environments arises from multiple agents acting concurrently, often following independent local heuristics (policies). This interaction leads to inherent uncertainty in the outcomes (returns), especially during the initial phases of training. Although exploration strategies can help reduce the episodic uncertainty by utilizing the generalization capability of neural networks, it remains an integral part of multi-agent environments at the start of training [14, 18]. To address this challenge, we propose leveraging the safety implications available from the environment, e.g., agent deaths, to boost convergence during policy learning. We use a loss function based on Control Barrier Function (CBF) [22, 41], alongside the Huber quantile loss, for training the local policies. In the following sub-section we formally describe the Barrier Function loss used in this work.

### 4.1 The Barrier Function Loss

Barrier Function is employed to ensure that the agents' trajectories stay within a "safe" region of the state space. This barrier defines constraints that the policies must satisfy, ensuring that unsafe states are avoided during training. Formally, the barrier certificate is defined as follows [41].

**Definition 1. Barrier Certificate.** A barrier certificate,  $B^\pi$ , for a policy,  $\pi$  is a function  $B^\pi : \mathcal{S} \rightarrow \mathbb{R}$ , such that the following properties hold.

- (1)  $\forall s \in \mathcal{S}_0, B^\pi(s) \leq 0$
- (2)  $\forall s \in \mathcal{S}_{\text{unsafe}}, B^\pi(s) > 0$
- (3)  $\forall s \in \mathcal{S}, B^\pi(s') - B^\pi(s) \leq -\lambda_B B^\pi(s)$

where  $s' \sim P(\cdot|s, \mathbf{u})$ ,  $\mathbf{u} \sim \pi$ , and  $0 < \lambda_B < 1$  is a hyper-parameter controlling the convergence rate.

In our work, the barrier function is designed to capture the number of agent terminations as a function of the state over the course of a trajectory. Specifically, the barrier function is defined using the following equation.

$$B^\pi(s) = (\text{agents dead at } s) + \gamma_B B^\pi(s') \quad (4)$$

where  $s' \sim P(\cdot|s, \mathbf{u})$ , and  $\gamma_B$  signifies the discount factor for the barrier function. Intuitively, it can be thought of as an estimation of vulnerability (opposite of agent health) of the collective group of agents. The formulation bears similarity with the definition of the value function  $V_\pi$ , as it includes the discounted future state estimate at  $s'$  in the current estimate; since in an MDP the state  $s' \sim P(\cdot|s, \cdot)$ . Among the three conditions above, our policy should satisfy the third condition (which is the invariant property), i.e.,  $\forall s \in \mathcal{S}, B^\pi(s') - B^\pi(s) \leq -\lambda_B B^\pi(s)$ , ensuring that the barrier function consistently decreases along the trajectory [41].

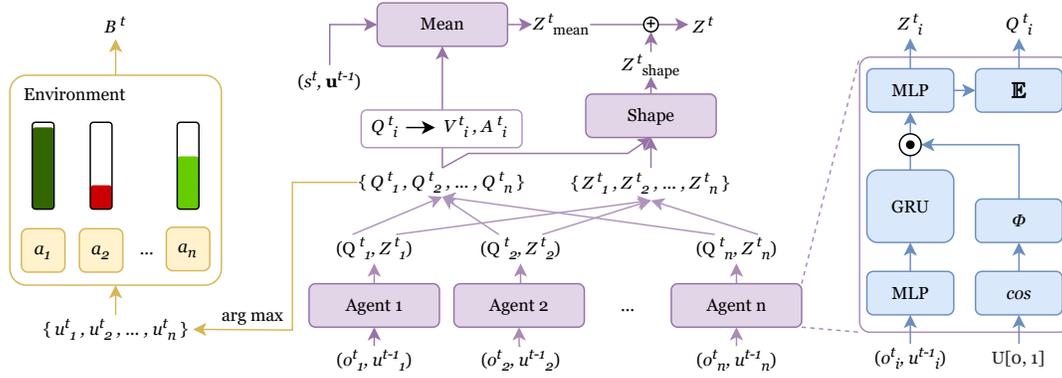
To enforce this condition, we integrate the following loss function for policy optimization:

$$\mathcal{L}_B(\pi) = \frac{1}{|\mathcal{S}|} \sum_{\substack{s, s' \in \mathcal{S} \\ s' \sim P(\cdot|s, \mathbf{u})}} \max(B^\pi(s') - (1 - \lambda_B)B^\pi(s), 0) \quad (5)$$

where  $|\mathcal{S}|$  represents the number of visited states over a trajectory. This loss function is derived from the number of agent deaths, which is a safety-critical metric. By focusing on this deterministic safety criterion, the method improves the accuracy of policy updates in unsafe scenarios.

An alternative approach could be to calculate the barrier loss for each individual agent, and update the local policies  $\pi_i$  accordingly. But this approach has two notable drawbacks; • given the length of the trajectory, the impact of this loss at each state may be insufficient to derive meaningful gradient updates, particularly when averaged across the trajectory, • in a cooperative environment, it is unrealistic to attribute the responsibility of an agent's termination solely to the corresponding local policy, as the outcome is influenced by the actions of all agents.

Thus, our method leverages the global safety constraint across the trajectory, promoting stable and efficient convergence in MARL.



**Figure 2: The overall framework illustrating the training objective. The calculated errors are combined using gradient manipulation and back-propagated through the central training model. References to the abbreviated notations from the text:  $Z_i^t = Z_{\pi_i}(o_i^t, u_i^t)$ ,  $Q_i^t = Q_{\pi_i}(o_i^t, u_i^t)$ ,  $t$  is time,  $A_i^t = A_{\pi_i}(o_i^t, u_i^t)$ ,  $V_i^t = V_{\pi_i}(o_i^t)$ .**

## 4.2 Optimization on Return

After defining the safety objective through the CBF loss, we integrate it with the DRL objective to optimize the reward function. We use the Huber-quantile regression loss [8] for  $Q_{\pi}$ , allowing us to model the stochastic return distribution effectively.

In cooperative multi-agent systems, value factorization techniques assist CTDE. However, to account for the stochasticity of value functions, we require a distributional variant of value factorization that satisfies Distributional IGM [33]. The Mean-Shape decomposition [33, 34] addresses this by separating distribution  $Z$  over global utilities into:  $Z = \mathbb{E}[Z] + (Z - \mathbb{E}[Z]) = Z_{\text{mean}} + Z_{\text{shape}}$ . The mean term in the mean-shape decomposition  $\mathbb{E}[Z]$ , represents the expected value, while the shape term captures deviations from this expectation [17, 34]. This ensures compliance with DIGM while allowing effective factorization of stochastic utilities.

For factorization, we adopt a distributional version of the QPLEX algorithm [38], using a Dueling Network Architecture [39]. Following [35], the global action-value function  $Q_{\pi}$  is factorized as follows.

$$Q_{\pi}(s, \mathbf{u}) = \sum_{i=1}^n V_{\pi_i}(o_i) + \sum_{i=1}^n \lambda_i(o_i, u_i) \cdot A_{\pi_i}(o_i, u_i) + \sum_{i=1}^n (Z_{\pi_i}(o_i, u_i) - Q_{\pi_i}(o_i, u_i)) \quad (6)$$

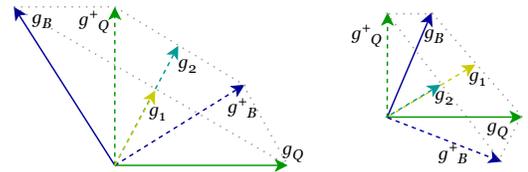
Here  $V_{\pi_i}(o_i)$  denotes the value function for agent  $i$ 's observation,  $A_{\pi_i}(o_i, u_i)$  is the advantage function, and  $Z_{\pi_i}(o_i, u_i)$  captures the distributional nature of local utilities. The overall architecture for obtaining the global distribution over returns and the CBF value from the environment is illustrated in Figure 2.

## 4.3 Overall Training Objective

Now that we have obtained the formulations for calculating both the loss functions, we need to integrate them to calculate our overall loss for the gradient descent. In the following formulation for calculating the overall loss, we consider the Huber quantile loss as  $\mathcal{L}_Q$  and the CBF loss as  $\mathcal{L}_B$ . We combine the loss functions using

gradient manipulation, where the gradient is calculated from the back-propagation of the individual loss function components, as studied in [42] for multi-task learning.

The objective of gradient manipulation is to reduce the deviation of gradients due to the reward and barrier components. Let the gradient due to  $\mathcal{L}_B$  be represented as  $g_B$ , and that due to the  $\mathcal{L}_Q$  as  $g_Q$ . Let  $\theta$  be the angle between  $g_B$  and  $g_Q$ . Since we work with gradient manipulation here, the gradients occasionally may need to be projected to a plane that is normal to the other gradient; we will explain this in more detail in the following paragraph. Let the projection of gradient of  $g_B$  on the normal plane of  $g_Q$  be given by  $g_B^+$ ; while the projection of gradient of  $g_Q$  on the normal plane of  $g_B$  be represented by  $g_Q^+$ .



**Figure 3: Analysis of gradient manipulation.**

Owing to optimization on two tasks, there may occasionally occur a conflict among the gradients  $g_Q$  and  $g_B$ , i.e.  $\theta > 90^\circ$ . In those cases, it is essential to project the gradient vector onto the normal plane of the other vector for effective updates to the network [42]. For cases when  $\theta \leq 90^\circ$ , i.e., the gradients do not conflict, the gradient manipulation is not required. The overall gradient calculation,  $g$ , is thus given by the following equation.

$$g = \begin{cases} \beta_Q^+ \cdot g_Q^+ + \beta_B^+ \cdot g_B^+, & \theta > 90^\circ = g_1(\text{let}) \\ \beta_Q \cdot g_Q + \beta_B \cdot g_B, & \theta \leq 90^\circ = g_2(\text{let}) \end{cases} \quad (7)$$

$$g_Q^+ = g_Q - \frac{g_Q \cdot g_B}{\|g_B\|^2} g_B \quad g_B^+ = g_B - \frac{g_B \cdot g_Q}{\|g_Q\|^2} g_Q \quad (8)$$

where  $\beta_Q^+$ ,  $\beta_B^+$ ,  $\beta_Q$ ,  $\beta_B$  denote the weights of the gradients. With this strategy we can optimize on both the tasks more efficiently. In

Equation 7,  $g_1$  is the projection of gradient onto the normal plane of the other gradient. As shown in Equation 7, we have two cases for the gradients corresponding to the loss functions calculated: when the gradients conflict, i.e.,  $\theta > 90^\circ$ ; and when the gradients do not conflict, i.e.,  $\theta \leq 90^\circ$ . The process is illustrated in Figure 3.

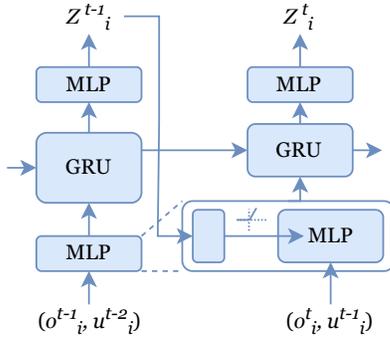
1. *When  $\theta > 90^\circ$ :* We consider  $\beta_Q^+$  and  $\beta_B^+$  are equal to 0.5. Then according to the given equations, we have the following.

$$\begin{aligned} g_1 &= 0.5 \cdot g_Q^+ + 0.5 \cdot g_B^+ \\ &= 0.5 \cdot \left( g_Q - \frac{g_Q \cdot g_B}{\|g_B\|^2} g_B \right) + 0.5 \cdot \left( g_B - \frac{g_B \cdot g_Q}{\|g_Q\|^2} g_Q \right) \\ g_2 &= 0.5 \cdot g_Q + 0.5 \cdot g_B \end{aligned}$$

We know that, when  $\theta > 90^\circ$ ,  $(g_Q \cdot g_B) < 0 \Rightarrow -(g_Q \cdot g_B) > 0$ . Therefore, according to the above equations,  $g_1 > g_2$ . So when  $\theta > 90^\circ$ , we use  $g_1$  as our gradient.

2. *When  $\theta \leq 90^\circ$ :* We consider  $\beta_Q$  and  $\beta_B$  are equal to 0.5. We know that, when  $\theta \leq 90^\circ$ ,  $(g_Q \cdot g_B) \geq 0 \Rightarrow -(g_Q \cdot g_B) \leq 0$ . Therefore, according to the above equations,  $g_1 \leq g_2$ . So when  $\theta \leq 90^\circ$ , we use  $g_2$  as our gradient.

#### 4.4 The Local Policy Network



**Figure 4: The input layer weights of  $\pi_i$  generated using a hyper-network from previous step’s return distribution.**

Now we focus on the architecture of the local policy,  $\pi_i$ . As illustrated in Figure 2, we adopt IQN [8] for sampling distributions over local actions. To address partial observability in local observations, we use a recurrent neural network, as recommended by [11].

The input layer of the local policy is designed using a hyper-network layer, as shown in Figure 4. Since  $\pi_i$  predicts a distribution over returns, representing future rewards, we leverage it to dynamically predict weights for the input layer. The hyper-network thus allows for more adaptable weight generation based on the distributional information from future states. A ReLU activation is used to filter the non-negative weights motivated by two conditions: • negative influences on individual observation elements can propagate through subsequent layers, potentially distorting the representation of other elements, • given the inherent uncertainty in the return, an optimistic approach with non-negative weights is preferable. For clarity, Figure 4 omits the details of IQN, focusing instead on the structure of the hyper-network.

## 5 THEORETICAL ANALYSIS

In this section we study theoretical guarantees on convergence using our approach and provide sample complexity advancements. Since we are not proposing a new methodology for value decomposition here, we therefore assume the convergence of the background value decomposition methodology used for MARL, studied in more detail here [37]. Additionally, we assume Policy Gradient as our optimization approach; and for proving convergence, we assume a parameterized policy  $\pi_m$  (with policy parameters  $m$ ) in tabular setting with direct policy parameterization.

**THEOREM 5.1.** *Consider a tabular setting for our policy with direct policy parameterization, with the number of policy parameters,  $m \in \Delta(\mathcal{U})^{|\mathcal{S}|}$ . Let the step size for the NPG update be  $\alpha = (1 - \gamma)^{1.5} / \sqrt{|\mathcal{S}||\mathcal{U}|T}$ , where  $T$  is the number of policy update epochs. Then, with confidence  $(1 - \delta)$ , and assuming convergence of the underlying TD update algorithm, we have the following.*

$$V_{\pi^*}(s_0) - \mathbb{E}[V_{\pi_{m_T}}(s_0)] \leq \Theta \left( \sqrt{\frac{|\mathcal{S}||\mathcal{U}|}{(1 - \gamma)^3 T}} \right)$$

The proof for the above theorem is provided in the Supplementary Material<sup>1</sup>. In the following sub-section we formally define our safety criterion, and provide a probabilistic confidence bound on the safety violation for the defined safety criterion.

### 5.1 Safety Verification

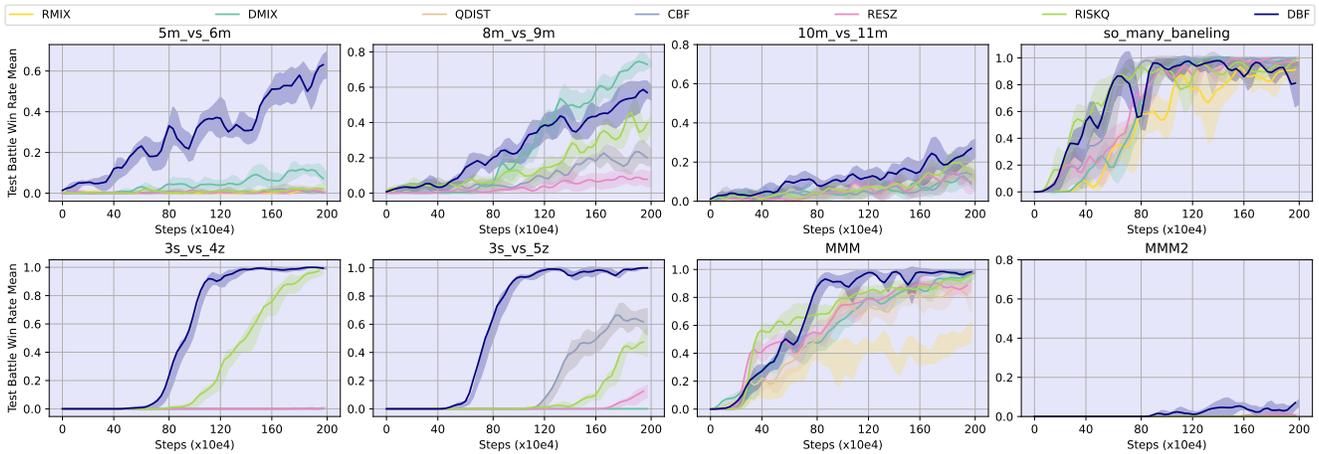
With the safety objective, we can verify the safety of the system by introducing a probabilistic verification method for the system. In (MA)RL, classical methods for safety verification that take binary decision about whether the system is safe or unsafe, cannot be used, since the system dynamics or the knowledge about the unsafe states are not available. So we address the safety verification from a quantitative safety perspective; i.e., we determine the probability of obtaining an unsafe trajectory remaining within a safety target with some confidence. Here a safe trajectory means reaching the goal state with some agents still alive.

Our barrier function, although calculated from agent terminations, is not a direct indicator of the terminations. However, we would like to express our constraint in terms of agent terminations. So let  $V_{\pi}^B(\tau)$  be the sum of agent terminations over the trajectory, where  $\tau$  is a test trajectory sampled with the trained policy  $\pi$ . Therefore, we can write our objective as a chance-constrained optimization program (CCP) [6] as: given a parameter  $\beta \in (0, 1)$ , our objective is to compute  $\varepsilon \in (0, 1)$  such that with confidence  $\geq (1 - \beta)$  our safety constraint is satisfied, i.e.,  $\tau$  is safe, with probability  $\geq (1 - \varepsilon)$ . Mathematically,  $\tau$  is considered safe if  $V_{\pi}^B(\tau)$  is below a given safety threshold  $\omega$ . Our objective can be formally restated as, with confidence at least,  $(1 - \beta)$ ,

$$P(V_{\pi}^B(\tau) \leq \omega) \geq (1 - \varepsilon) \tag{9}$$

Note that  $\omega$  can be considered equal to 0; but it may have some disadvantages. When it is evident that there may be casualties in the path to the goal state, taking  $\omega = 0$  may harm the return since the policy will deviate from the optimal possible policy owing to the gradients from the barrier function loss. Thus  $\omega$  is kept to be

<sup>1</sup> The full paper with Supplementary Material can be found [here](#).



**Figure 5: Comparison with the state-of-the-art algorithms using the StarCraft hard and super-hard scenarios. Our method is marked as DBF (Distributional MARL with a Barrier Function based on agent casualties).**

some positive number dependent on the number of agents,  $n$ ; e.g.,  $\omega = (n - 1)$ . Finally, from equations 1 and 9, we can rewrite our constrained optimization problem as follows.

$$\begin{aligned} & \arg \max_{\pi} \mathbb{E} \left[ \sum_{t \sim \tau} \gamma^t r(s_t, \mathbf{u}_t) \right] \\ \text{s.t. } & P(V_{\pi}^B(\tau) \leq \omega) \geq (1 - \epsilon) \end{aligned}$$

We can solve the above CCP problem using  $\epsilon$  a sampling approach in place of the probabilistic constraint, where we substitute the constraint with  $N$  i.i.d. sample constraints as:  $V_{\pi}^B(\tau_i) \leq \omega, \forall i \in [1, N]^2$ . The resulting solution may not be a good approximation of the CCP solution. So for improving the robustness of the solution, we remove some  $k$  samples from  $N$  [6, 29]. In other words, we can say that we are removing  $k$  out of  $N$  constraints from our original problem. The number of removed constraints has a direct impact on the safety probability  $(1 - \epsilon)$ . The sample-based optimization of the above problem can therefore be expressed as follows.

$$\begin{aligned} & \arg \max_{\pi} \mathbb{E} \left[ \sum_{t \sim \tau} \gamma^t r(s_t, \mathbf{u}_t) \right] \\ \text{s.t. } & V_{\pi}^B(\tau_i) \leq \omega, \forall i \in \{1, \dots, N\} - \mathcal{A}(\{1, \dots, N\}) \end{aligned}$$

where  $\mathcal{A}$  is some selection rule. Let  $P^N$  be the product of probabilities for  $N$  samples. Using the theorem below from [6, 29], we can establish the relation between  $\epsilon$  and the confidence parameter,  $\beta$ , given the values of  $N$  and  $k$ .

**THEOREM 5.2 ([29]).** *Given a small confidence parameter  $\beta \in (0, 1)$ , if  $N$  and  $k$  are such that,*

$$\binom{k+m-1}{k} \sum_{i=0}^{k+m-1} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i} \leq \beta$$

where  $m$  is the number of policy parameters, then  $P^N(P(V_{\pi}^B(\tau_i) > \omega) \leq \epsilon) \geq (1 - \beta)$ .

<sup>2</sup>Please note that here  $N$  represents  $N$  i.i.d. samples, which is different from  $N$  that represents the set of agents.

Using the above theorem and from Equation 8 in [6] we know,

$$\begin{aligned} k & \leq \epsilon N - m + 1 - \sqrt{2\epsilon N \ln \frac{(\epsilon N)^{m-1}}{\beta}} \\ \Rightarrow \frac{(\epsilon N)^{m-1}}{\beta} & \leq e^{(k+m-1-\epsilon N)^2 / (2\epsilon N)} \end{aligned}$$

Let us consider  $(k + m - 1 - \epsilon N)^2 / (2\epsilon N) = \psi$ . Therefore;

$$\epsilon \leq \left( e^{\ln \beta / (m-1)} \times e^{\psi / (m-1)} \right) / N$$

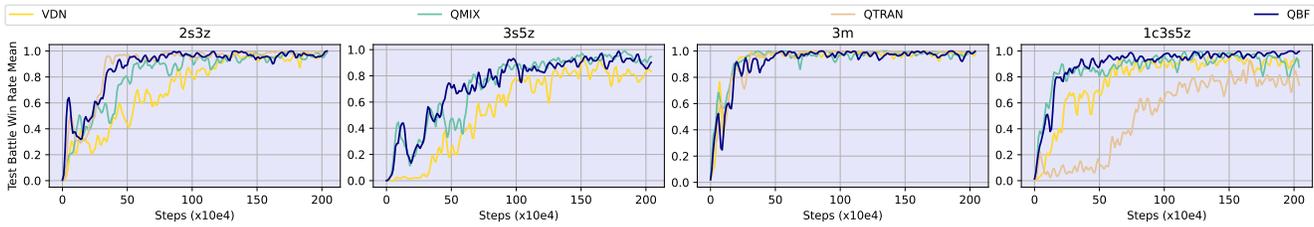
This is an implicit inequality with  $\epsilon$  occurring on both sides. However, we can deduce that, if we consider  $m$  to be finite quantity less than  $N$ , for small value of  $\beta$  and  $k$ , the right hand side of the above equation will be very small. Thus  $\epsilon$ , which signifies the upper bound on the probability of unsafe trajectories, will be very small.

## 6 EMPIRICAL STUDIES

In this section we present empirical evaluations of our method. The primary objective is not to introduce a novel constrained MARL approach but to evaluate the impact of incorporating agent casualties as an implicit constraint into the optimization process. The source code is available in GitHub. Our approach is compared with state-of-the-art Distributional MARL algorithms on the StarCraft II [26] and MetaDrive [15] benchmark environments.

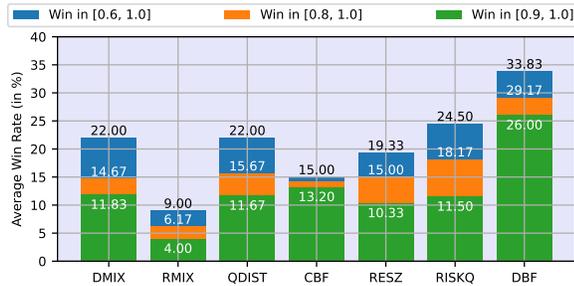
### 6.1 StarCraft Battle Scenarios

The StarCraft Multi-Agent Control scenarios include various battle scenarios where allies, controlled by our MARL algorithm, face enemy units controlled by built-in rules. The objective is to maximize the win rate while minimizing ally casualties [26]. We used the PyMARL framework for our experimental setup and training our algorithms. The evaluation metric is the fraction of number of battles won during testing. We derive our barrier function based on the objective of ally casualties and integrate with the reward optimization algorithm to align with the task's objectives. The action space for each of the agent is discrete, thus allowing us to



**Figure 6: Comparison with the baselines using the StarCraft easy scenarios. Our method is marked as QBF (Q-function evaluation in MARL with a Barrier Function based on agent casualties).**

derive distributions over the returns corresponding to each action. We evaluated our approach on two fronts: • firstly we compare with the distributional MARL approaches on hard and super-hard scenarios; • secondly we compare with MARL approaches on easy scenarios. For the second case, we do not have distributions over the returns, and cannot use them for our policy network as proposed in sub-section 4.4.

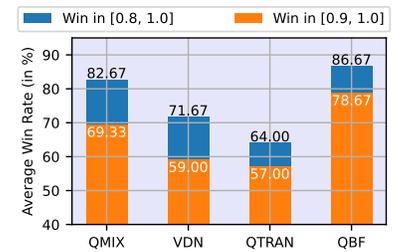


**Figure 7: Summary of the evaluations on the StarCraft hard and super-hard scenarios.**

For the first case, we used the following distributional MARL algorithms for comparison: • RMIX [23] uses CVaR with QMIX • DMIX [33] extending QMIX to the distributional setting • QDIST [31] separately analyses the risk measures from the return distributions, based on source of error • CBF [22] integrates with DMIX an action correction strategy using a decentralized barrier function approximated using a separate network • RESQ [28] uses residual Q-functions • RISKQ [27] uses local risk sensitive action selection policies. Our algorithm, referred to as DBF (Distributional MARL with Barrier Function), integrates the barrier loss with DMIX. We used on-policy samples for calculating our barrier function loss [41]. The results are shown in Figure 5. Our algorithm is marked as DBF, which is short for Distributional MARL with a Barrier Function based on agent casualties. Results in Figure 7 indicate DBF achieves a higher fraction of test wins across varying thresholds ( $\geq 0.6$ ,  $\geq 0.8$ ,  $\geq 0.9$ ). From the results, it can be inferred that our approach improves the accuracy of the distributional MARL approaches for the hard and super-hard scenarios. The hard and super-hard scenarios generally contain uneven distribution of agents between the ally and enemy teams; with the enemy team generally containing more number of agents than the ally team. This makes winning battle tougher in these scenarios. We used  $\gamma_B = 0.5$  for our experiments.

The hyper-parameters used for our experiments are summarized in the Supplementary Material<sup>1</sup>.

We also compared our algorithm with traditional MRL approaches: • VDN [36] which uses summation of local returns • QMIX [24] uses monotonic factorization of value function • QTRAN [30] uses transformers for factorization. The traditional MARL approaches has been demonstrated to perform well for the easy scenarios in StarCraft.



**Figure 8: Summary of evaluations on StarCraft easy scenarios.**

The easy scenarios contain equal number of ally and enemy agents with similar types of agents for both groups. The action space is discrete here. We integrated our barrier function loss with the QMIX algorithm for evaluation here. The results are shown in figures 6 and 8. Our algorithm is marked as QBF, which stands for Q-function evaluation in MARL with a Barrier Function based on agent casualties. From the evaluations it is evident that integrating the barrier loss improves the performance of the baseline algorithm, albeit marginally. This experiment also highlights the importance of the hyper-network used in the local policy network, explained in Section 4.4.

## 6.2 MetaDrive Simulator

MetaDrive [15] is a lightweight and realistic driving simulator featuring various multi-agent cooperative driving tasks with decentralized rewards. It presents the challenge of controlling multiple cars in some predefined scenarios. Since the simulator does not return any global state; for training certain algorithms, we used the collective observation vector as the global state. Each agent (vehicle) can execute discrete actions, with crashes or road exits resulting in termination. We did not allow re-spawn of vehicles for the terminated vehicles. Episodes were terminated when more than half the agents are eliminated during training; or when any agent reached termination during evaluation. We used 10 agents for each environment for training and evaluation.

For our experiments, we used a global reward computed by summing individual agent rewards to align with the POMDP structure.

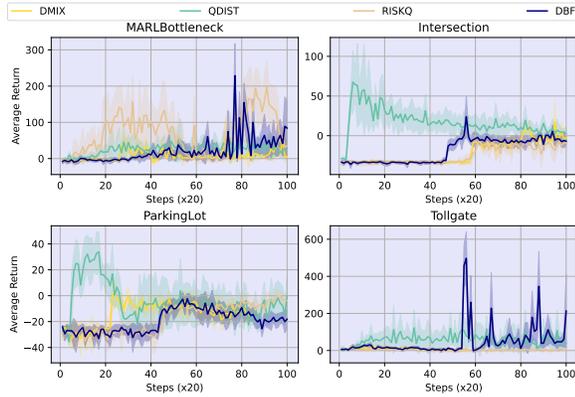


Figure 9: Evaluation results on the Metadrive MARL environments.

Training spanned 2000 episodes with a replay buffer holding the most recent 1000 episodes. On-policy samples were used for barrier loss calculation. The trained policy was evaluated on 10 episodes at each evaluation step. As shown in Fig. 9, our approach demonstrates strong performance compared to baseline algorithms across scenarios. Key hyper-parameters include  $\gamma = 0.99$ ,  $\gamma_B = 0.5$ , and the Adam optimizer. Additional experimental details are provided in the Supplementary Material<sup>1</sup>.

### 6.3 Ablation Studies

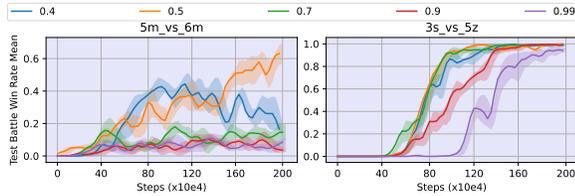


Figure 10: Ablation Study for  $\gamma_B$  tested on StarCraft.

In our experiments, for the discount factor of barrier function, we had used  $\gamma_B = 0.5$ . We had tested various values for the hyper-parameter on two StarCraft scenarios before using them for the rest of our evaluations. The hyper-parameter dictates the credit assignment horizon length based on the fault metric, i.e., number of agent terminations. From the results in Figure 10, it is evident that the algorithm performs well for  $\gamma_B = 0.5$ . The tested values for  $\gamma_B$  were 0.4, 0.5, 0.7, 0.9, 0.99.

We also conducted an Ablation Study on adjustment of weights of the gradients,  $\beta_Q$  and  $\beta_B$  on two StarCraft environments. The results are shown in Figure 11. Here we only show scores for  $\beta_Q$ , as  $\beta_B = 1 - \beta_Q$ . The bars in the figure indicate the percentage of evaluations for which the win-rate was  $\geq 0.6$ ,  $\geq 0.8$ , or  $\geq 0.9$ . From the results it is evident that setting equivalent weights for both the gradients, i.e.,  $\beta_Q = \beta_B = 0.5$ , results in better performance for most scenarios, rather than weighing more on any one of the gradients. We used this result for the rest of our evaluations.

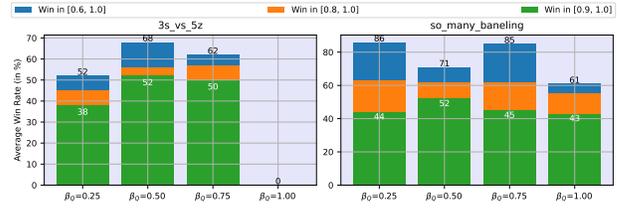


Figure 11: Ablation Study for  $\beta_B$  and  $\beta_Q$  on StarCraft.

## 7 CONCLUSION

In this work, we proposed a novel approach to distributional multi-agent reinforcement learning (MARL) that integrates fault tolerance based on an agent safety boundary (depending on intrinsic environment dynamics) with effective reward optimization. By incorporating a constraint-based formulation to model agent termination as a crucial learning signal, our method addresses the issue of unsafe actions leading to agent failures, thus improving the stability of MARL in environments where agent survivability is critical. We extended our approach for both traditional MARL value factorization techniques and distributional MARL approaches since the safety constraint is not limited to the distributional analysis. We also suggested an improvement in the local policy network to utilize the predicted distributions over returns.

The integration of these techniques using the PCGrad optimizer resulted in an effective hybrid loss formulation, combining the distributional Huber-quantile loss with our barrier function based safety loss. This allowed us to strike a balance between optimizing agent performance and ensuring system-wide safety. Our experimental results, particularly on complex environments such as StarCraft II micromanagement tasks, and MetaDrive multi-agent driving scenarios, demonstrate that our approach performs comparatively well as compared with the state-of-the-art MARL methods in terms of both convergence and policy robustness. We also theoretically analysed the convergence and safety verification of our approach. The approach can be further be integrated with other constrained MARL algorithms to test its robustness in environments with extrinsic safety constraints.

## ACKNOWLEDGMENTS

The authors acknowledge partial support from AI4CPS IIT Kharagpur grant no TRP3RDTR001 for this work.

## REFERENCES

- [1] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marl-book.com>
- [2] Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. 2018. Safe reinforcement learning via shielding. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [3] Osbert Bastani. 2021. Safe reinforcement learning with nonlinear dynamics via model predictive shielding. In *2021 American control conference (ACC)*. IEEE, 3488–3494.
- [4] Marc G Bellemare, Will Dabney, and Rémi Munos. 2017. A distributional perspective on reinforcement learning. In *International conference on machine learning*. PMLR, 449–458.
- [5] Suda Bharadwaj, Roderik Bloem, Rayna Dimitrova, Bettina Könighofer, and Ufuk Topcu. 2019. Synthesis of minimum-cost shields for multi-agent systems. In *2019 American Control Conference (ACC)*. IEEE, 1048–1055.

- [6] Marco C Campi and Simone Garatti. 2011. A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality. *Journal of optimization theory and applications* 148, 2 (2011), 257–280.
- [7] Yu-Han Chang, Tracey Ho, and Leslie Kaelbling. 2003. All learning is local: Multi-agent learning in global reward games. *Advances in neural information processing systems* 16 (2003).
- [8] Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. 2018. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning*. PMLR, 1096–1105.
- [9] Will Dabney, Mark Rowland, Marc Bellemare, and Rémi Munos. 2018. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [10] Ingy ElSayed-Aly, Suda Bharadwaj, Christopher Amato, Rüdiger Ehlers, Ufuk Topcu, and Lu Feng. 2021. Safe multi-agent reinforcement learning via shielding. *arXiv preprint arXiv:2101.11196* (2021).
- [11] Matthew Hausknecht and Peter Stone. 2015. Deep recurrent q-learning for partially observable mdp. In *2015 aaii fall symposium series*.
- [12] Wan Ju Kang, David Earl Hostallero, Kyunghwan Son, Daewoo Kim, and Yung Yi Qtran. 2019. Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *Proceedings of the 31st International Conference on Machine Learning, Proceedings of Machine Learning Research*. PMLR.
- [13] Jifeng Hu, Yanchao Sun, Hechang Chen, Sili Huang, Yi Chang, Lichao Sun, et al. 2022. Distributional reward estimation for effective multi-agent deep reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 12619–12632.
- [14] Yiding Jiang, J Zico Kolter, and Roberta Raileanu. 2024. On the importance of exploration for generalization in reinforcement learning. *Advances in Neural Information Processing Systems* 36 (2024).
- [15] Quanyi Li, Zhenghao Peng, Zhenghai Xue, Qihang Zhang, and Bolei Zhou. 2021. MetaDrive: Composing Diverse Driving Scenarios for Generalizable Reinforcement Learning. *arXiv preprint arXiv:2109.12674* (2021).
- [16] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* 30 (2017).
- [17] Clare Lyle, Marc G Bellemare, and Pablo Samuel Castro. 2019. A comparative analysis of expected and distributional reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4504–4511.
- [18] Borislav Mavrin, Hengshuai Yao, Linglong Kong, Kaiwen Wu, and Yaoliang Yu. 2019. Distributional reinforcement learning for efficient exploration. In *International conference on machine learning*. PMLR, 4424–4434.
- [19] Jihwan Oh, Joonkee Kim, Minchan Jeong, and Se-Young Yun. 2023. Toward risk-based optimistic exploration for cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2303.01768* (2023).
- [20] Jihwan Oh, Joonkee Kim, and Se-Young Yun. 2022. Risk perspective exploration in distributional reinforcement learning. *arXiv preprint arXiv:2206.14170* (2022).
- [21] Frans A Oliehoek, Matthijs TJ Spaan, and Nikos Vlassis. 2008. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research* 32 (2008), 289–353.
- [22] Zengyi Qin, Kaiqing Zhang, Yuxiao Chen, Jingkai Chen, and Chuchu Fan. 2021. Learning safe multi-agent control with decentralized neural barrier certificates. *arXiv preprint arXiv:2101.05436* (2021).
- [23] Wei Qiu, Xinrun Wang, Runsheng Yu, Xu He, Rundong Wang, Bo An, Svetlana Obraztsova, and Zinovi Rabinovich. 2020. RMIX: Risk-sensitive multi-agent reinforcement learning. (2020).
- [24] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research* 21, 178 (2020), 1–51.
- [25] Mark Rowland, Robert Dadashi, Saurabh Kumar, Rémi Munos, Marc G Bellemare, and Will Dabney. 2019. Statistics and samples in distributional reinforcement learning. In *International Conference on Machine Learning*. PMLR, 5528–5536.
- [26] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder De Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043* (2019).
- [27] Siqi Shen, Chennan Ma, Chao Li, Weiquan Liu, Yongquan Fu, Songzhu Mei, Xinwang Liu, and Cheng Wang. 2024. RiskQ: risk-sensitive multi-agent reinforcement learning value factorization. *Advances in Neural Information Processing Systems* 36 (2024).
- [28] Siqi Shen, Mengwei Qiu, Jun Liu, Weiquan Liu, Yongquan Fu, Xinwang Liu, and Cheng Wang. 2022. Resq: A residual q function-based approach for multi-agent reinforcement learning value factorization. *Advances in Neural Information Processing Systems* 35 (2022), 5471–5483.
- [29] Arambam James Singh and Arvind Easwaran. 2024. PAS: Probably Approximate Safety Verification of Reinforcement Learning Policy Using Scenario Optimization. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 1745–1753.
- [30] Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International conference on machine learning*. PMLR, 5887–5896.
- [31] Kyunghwan Son, Junsu Kim, Sungsoo Ahn, Roben D Delos Reyes, Yung Yi, and Jinwoo Shin. 2022. Disentangling sources of risk for distributional multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 20347–20368.
- [32] Kyunghwan Son, Junsu Kim, Yung Yi, and Jinwoo Shin. 2021. Disentangling sources of risk for distributional multi-agent reinforcement learning. (2021).
- [33] Wei-Fang Sun, Cheng-Kuang Lee, and Chun-Yi Lee. 2021. DFAC framework: Factorizing the value function via quantile mixture for multi-agent distributional q-learning. In *International Conference on Machine Learning*. PMLR, 9945–9954.
- [34] Wei-Fang Sun, Cheng-Kuang Lee, and Chun-Yi Lee. 2021. A Distributional Perspective on Value Function Factorization Methods for Multi-Agent Reinforcement Learning. In *AAMAS*. 1671–1673.
- [35] Wei-Fang Sun, Cheng-Kuang Lee, Simon See, and Chun-Yi Lee. 2023. A unified framework for factorizing distributional value functions for multi-agent reinforcement learning. *Journal of Machine Learning Research* 24, 220 (2023), 1–32.
- [36] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 2017. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296* (2017).
- [37] Jianhao Wang, Zhizhou Ren, Beining Han, Jianing Ye, and Chongjie Zhang. 2021. Towards understanding cooperative multi-agent q-learning with value factorization. *Advances in Neural Information Processing Systems* 34 (2021), 29142–29155.
- [38] Jianhao Wang, Zhizhou Ren, Terry Liu, Yang Yu, and Chongjie Zhang. 2020. Qplex: Duplex dueling multi-agent q-learning. *arXiv preprint arXiv:2008.01062* (2020).
- [39] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. 2016. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*. PMLR, 1995–2003.
- [40] Yaodong Yang, Jianye Hao, Ben Liao, Kun Shao, Guangyong Chen, Wulong Liu, and Hongyao Tang. 2020. Qatten: A general framework for cooperative multiagent reinforcement learning. *arXiv preprint arXiv:2002.03939* (2020).
- [41] Yujie Yang, Yuxuan Jiang, Yichen Liu, Jianyu Chen, and Shengbo Eben Li. 2023. Model-free safe reinforcement learning through neural barrier certificate. *IEEE Robotics and Automation Letters* 8, 3 (2023), 1295–1302.
- [42] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems* 33 (2020), 5824–5836.
- [43] Wenbo Zhang, Osbert Bastani, and Vijay Kumar. 2019. Mamps: Safe multi-agent reinforcement learning via model predictive shielding. *arXiv preprint arXiv:1910.12639* (2019).