

Collaborative Adaptation for Recovery from Unforeseen Malfunctions in Discrete and Continuous MARL Domains

Yasin Findik, Hunter Hasenfus, and Reza Azadeh

Abstract—Cooperative multi-agent learning plays a crucial role for developing effective strategies to achieve individual or shared objectives in multi-agent teams. In real-world settings, agents may face unexpected failures, such as a robot’s leg malfunctioning or a teammate’s battery running out. These malfunctions decrease the team’s ability to accomplish assigned task(s), especially if they occur after the learning algorithms have already converged onto a collaborative strategy. Current leading approaches in Multi-Agent Reinforcement Learning (MARL) often recover slowly – if at all – from such malfunctions. To overcome this limitation, we present the Collaborative Adaptation (CA) framework, highlighting its unique capability to operate in both continuous and discrete domains. Our framework enhances the adaptability of agents to unexpected failures by integrating inter-agent relationships into their learning processes, thereby accelerating the recovery from malfunctions. We evaluated our framework’s performance through experiments in both discrete and continuous environments. Empirical results reveal that in scenarios involving unforeseen malfunction, although state-of-the-art algorithms often converge on sub-optimal solutions, the proposed CA framework mitigates and recovers more effectively.

I. INTRODUCTION

The popularity of multi-agent learning methods is rising due to their ability to manage the complexities of robotic systems and their importance in tasks that require detailed coordination [1]. These methods are valuable across various fields, such as robotics for search & rescue operations [2] and autonomous driving [3], whether involving multiple or single robot systems. In single-robot applications, multi-agent learning facilitates coordination among different system modules, exemplified by treating each joint of a robot manipulator as an individual agent within a collaborative team. In both multi-robot and single-robot systems, the collaborative behavior of the agents is particularly crucial for quickly and autonomously recovering from unexpected malfunctions, such as joint failures or battery failure that prevents the robot(s) from moving. In these situations, agents must improve their collaboration further and modify their joint strategies to successfully tackle these challenges.

In the field of Multi-Agent Reinforcement Learning (MARL) for cooperative tasks, the Centralized Training with Decentralized Execution (CTDE) [4] paradigm has emerged as a prominent approach, notably through recent advancements in deep reinforcement learning [5], [6], [7]. It effectively addresses a range of cooperative challenges in MARL, including the curse of dimensionality [8], non-stationarity [9]

Authors are with the Persistent Autonomy and Robot Learning (PeARL) Lab, Richard Miner School of Computer and Information Sciences, University of Massachusetts Lowell, MA, USA
yasin.findik@student.uml.edu, reza@cs.uml.edu

and global exploration [10]. However, despite its effectiveness in collaborative tasks, CTDE methods struggle with rapid adaptation to unforeseen agent malfunction(s) or failure. This challenge might be due to two primary reasons: the lack of built-in mechanisms to manage unforeseen robot failures, or the absence of design elements to steer agents’ collaboration strategies. Consequently, these algorithms face delays in adapting, as they need to independently discover and converge on new *effective* collaboration strategies without any guidance. To assist agents in searching for these strategies, we associate the discovery of such collaborative behaviors through inter-agent relationships, thereby accelerating and improving the adaptability of the agents as a team.

In this paper, we introduce a novel MARL framework that integrates a relational network with any multi-agent learning algorithm following the CTDE paradigm. This integration highlights the relative importance among agents, thereby enhancing their collaboration towards a specific strategy. This innovation facilitates faster and more effective adaptation to unexpected robot failures. We conducted experiments in two settings: a multi-agent discrete environment with multiple robots and a continuous environment [11] simulating an ant with multiple agents (an agent per leg). We demonstrate the effectiveness of our approach by comparing it with Independent Deep Q-Networks (IDQN) [12] and Value Decomposition Networks (VDN) [13] in the discrete environment, and with Independent Q-Functionals (IQF) [14] and Multi-Agent Deep Deterministic Policy Gradient (MADDPG) [15] in the continuous environment. Empirical results indicate that our method improves teamwork among agents, whether in configurations like a single robot with multiple legs/joints or multiple robots. More importantly, it enhances the ability to adapt to unexpected failure by utilizing relational networks.

II. BACKGROUND AND RELATED WORK

Reinforcement learning (RL) aims at enabling an agent to maximize its cumulative future rewards through learning from its interactions within an environment. This process is often represented as a Decentralized Markov Decision Process (DEC-MDP), characterized by a tuple $\langle S, \mathcal{A}, \mathcal{R}, \mathcal{T}, \gamma, N \rangle$. S , \mathcal{A} and \mathcal{R} indicate the true state of the environment, the joint set of agents’ action and reward, respectively. $\mathcal{T}(s, \mathcal{A}, s') : S \times \mathcal{A} \times S \mapsto [0, 1]$ is the dynamics function, defining the transition probability, where s and s' are the current and next state, respectively. $\gamma \in [0, 1)$ is the discount factor, and N represents the number of agents.

Q-Learning & Q-Functionals. Q-Learning [16] and Deep

Q-Networks (DQN) [7] are two prominent *value-based* approaches in single-agent RL with discrete action spaces. Q-Learning employs an action-value *table*, Q^π , for a policy π , defined as $Q^\pi(s, a) = \mathbb{E}[G | s^t = s, a^t = a]$, where G and t denote the return and time-step, respectively. This can be re-formulated recursively as:

$$Q^\pi(s, a) = \mathbb{E}_{s'}[R(s, a) + \gamma \mathbb{E}_{a' \sim \pi}[Q^\pi(s', a')]],$$

where R represents the reward function. However, DQN learns an action-value *function*, denoted as \hat{Q} , corresponding to the optimal policy by minimizing this loss:

$$L(w_{\text{PN}}) = \mathbb{E}_{s, a, r, s'}[(r + \gamma \max_{a'}(\hat{Q}(s', a'; w_{\text{TN}})) - \hat{Q}(s, a; w_{\text{PN}}))^2]. \quad (1)$$

The parameters for the prediction network are denoted by w_{PN} , while w_{TN} represents the parameters for the target network. w_{TN} are periodically updated with w_{PN} to enhance the stability of the learning process. DQN also utilizes an experience replay memory that hold tuples $\langle s, a, r, s' \rangle$, to enhance the learning stability further.

In these approaches, state-action pairs are directly associated with \mathcal{R} , as denoted by:

$$\hat{Q}(s, a) : (\mathcal{S} \times \mathcal{A}) \mapsto \mathcal{R}. \quad (2)$$

This direct mapping restricts their utility to scenarios within discrete environments (i.e., a finite set of actions) because it is unfeasible to calculate the action-value for every state-action pair in continuous domains. To address this limitation, Q-Functionals (QF) [17], restructure the mapping by decoupling the state and action elements, illustrated as:

$$\hat{Q}^{\text{F}}(s, a) : \mathcal{S} \mapsto (\mathcal{A} \mapsto \mathcal{R}). \quad (3)$$

Here, the state is initially mapped to a function that subsequently associates actions with \mathcal{R} . Essentially, each state is depicted as a function formed through the learning of basis function coefficients within the action domain. These state-representative functions enable the quick assessment of many actions via matrix operations between the action representations and the learned coefficients. Thus, QF adeptly manage continuous domains in single-agent contexts.

Within multi-agent systems, the basic application of these value-based approaches involves each agent i independently optimizing its respective action-value table or function, such as Independent Q-Learning (IQL) [18], Independent Deep Q-Networks (IDQN) [7], Independent Q-Functionals (IQF) [14]. These methods, however, encounter difficulties because the independent policy modifications by each agent make the environment appear non-stationary from their respective perspectives. This non-stationarity violates the stationary requirements of the Markov property, crucial for the stable convergence of these algorithms. As a simple alternative to independent learning, fully centralized learning [19] can be adopted, in which a singular controller is utilized across all agents, enabling the joint learning of a value function. Yet, this centralized strategy is characterized by computational demand and scalability problems, as the observation and action spaces grow exponentially with the number of agents, possibly resulting in intractability.

Value Function Factorization. In the context of joint action-value function learning, where complexity escalates exponentially with the number of agents, value function factorization methods emerge as a promising solution. Following the centralized training with decentralized learning (CTDE) [4] paradigm, these methods facilitate autonomous action execution by individual agents, while centralizing the integration of their strategies. Consequently, these methods resolve the non-stationarity issue inherent in independent learning through centralized training and tackle the scalability challenges associated with centralized learning by enabling decentralized execution.

VDN [13] and QMIX [20] stand out as two significant learning algorithms in action-value function factorization. These algorithms are characterized by using separate action-value function, denoted as \hat{Q}_i , for each agent $i \in \{1, \dots, N\}$. The primary distinction between these methods is their strategy for calculating the total action value, Q_{tot} . VDN achieves Q_{tot} by summing the Q_i s, as expressed in:

$$\hat{Q}_{\text{tot}} = \sum_{i=1}^N \hat{Q}_i(s, a_i), \quad (4)$$

whereas QMIX employs a state-dependent continuous monotonic function to integrate these values, denoted by:

$$\hat{Q}_{\text{tot}} := f_s(\hat{Q}_1(s, a_1), \dots, \hat{Q}_n(s, a_n)), \quad (5)$$

with the condition of $\frac{\partial f_s}{\partial \hat{Q}_i} \geq 0, \forall i \in \{1, \dots, N\}$. Often, these factorization algorithms use DQN for approximating the action-value function and aim to minimize the loss (1) in a centralized way, as follows:

$$L(w_{\text{PN}}) = \mathbb{E}_{s, u, r_{\text{team}}, s'}[(r_{\text{team}} + \gamma \max_{u'}(\hat{Q}_{\text{tot}}(s', u'; w_{\text{TN}})) - \hat{Q}_{\text{tot}}(s, u; w_{\text{PN}}))^2]. \quad (6)$$

where r_{team} is obtained by uniformly summing the agents individual rewards, and u is the joint action of the agents.

The current value function factorization methods have demonstrated proficiency in creating stable collaborative within discrete environments. However, even with their advanced iterations [21], [22], these approaches encounter challenges in formulating effective policies within continuous environments. In such settings, the action space is delineated by continuous vectors, thereby encompassing an unbounded set of potential actions [23].

Policy Gradient (PG). In single-agent RL with continuous action spaces Policy Gradient (PG) approaches [24] are essential. These methods aim to directly optimize the policy parameters, symbolized as θ , to maximize the expected return. The principal strategy is to modify θ in the direction of the policy's gradient, formulated as:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{s \sim p^{\pi}, a \sim \pi_{\theta}}[\nabla_{\theta} \log \pi_{\theta}(a|s) Q^{\pi}(s, a)], \quad (7)$$

where p^{π} denotes the distribution of states under the policy π . Generally, PG approaches differ from each other in their methods for estimating Q^{π} .

Moreover, Deterministic Policy Gradient (DPG) [25] adapts the PG theorem and represents policy as $\mu_{\theta} : \mathcal{S} \mapsto \mathcal{A}$

with a vector of n parameters $\theta \in \mathbb{R}^n$. The formulation of policy gradient becomes as:

$$\nabla_{\theta} J(\mu_{\theta}) = \mathbb{E}_{s \sim p^{\mu}} [\nabla_{\theta} \mu_{\theta}(s) \nabla_a Q^{\mu}(s, a)|_{a=\mu_{\theta}(s)}]. \quad (8)$$

Also, Deep Deterministic Policy Gradient (DDPG) [26] approach extends DPG by employing deep neural networks to approximate both the policy μ and critic \hat{Q}^{μ} .

Multi-Agent DDPG (MADDPG). MADDPG [15] extends DDPG by adhering CTDE paradigm, facilitating its application in multi-agent settings. The policy of agents are represented as $\pi = [\pi_1, \dots, \pi_N]$ with parameters $\theta = [\theta_1, \dots, \theta_n]$ in MADDPG. The optimization of these policies is conducted through the gradient of the expected return, which for each agent $i \in \{1, \dots, N\}$ is formulated as:

$$\nabla_{\theta_i} J(\theta_i) = \mathbb{E}_{x \sim p^{\mu}, a_i \sim \pi_i} [\nabla_{\theta_i} \log \pi_i(a_i | o_i) \hat{Q}_i^{\pi}(\mathbf{o}, \mathbf{a})],$$

where $\hat{Q}_i^{\pi}(\mathbf{o}, \mathbf{a})$ denotes the centralized action-value function, incorporating both the states $\mathbf{o} = [o_1, \dots, o_N]$ and actions $\mathbf{a} = [a_1, \dots, a_N]$ of all agents to compute the Q -value. The framework is further adapted for deterministic policies $\mu = [\mu_{\theta_1}, \dots, \mu_{\theta_N}]$, modifying the gradient expression as:

$$\nabla_{\theta_i} J(\mu_{\theta_i}) = \mathbb{E}_{\mathbf{o} \sim p^{\mu}, \mathbf{a} \sim \mu} [\nabla_{\theta_i} \mu_{\theta_i}(a_i | o_i) \nabla_{a_i} \hat{Q}_i^{\mu}(\mathbf{o}, \mathbf{a})|_{a_i=\mu_{\theta_i}(o_i)}].$$

MADDPG, similar to the DQN algorithm, incorporates target networks and experience replay to enhance the stability of the policy networks. The loss function for updating the centralized action-value function Q_i^{μ} is defined as:

$$L(\theta_i) = \mathbb{E}_{\mathbf{o}, \mathbf{a}, \mathbf{r}, \mathbf{o}' \sim \mathcal{D}} [(r_i + \gamma Q_i^{\mu'}(\mathbf{o}', \mathbf{a}')|_{\mathbf{a}'=\mu'(\mathbf{o}')} - Q_i^{\mu}(\mathbf{o}, \mathbf{a}))^2],$$

where \mathcal{D} denotes the replay memory storing tuples $\langle \mathbf{o}, \mathbf{a}, \mathbf{r}, \mathbf{o}' \rangle$ and $\mu' = [\mu_{\theta'_1}, \dots, \mu_{\theta'_N}]$ represents the periodically updated target policies. Despite MADDPG's state-of-the-art performance in maximizing agents' return, it faces challenges like converging to local optima due to inefficient sampling and limited exploration in complex environments with extensive state and action spaces, specifically recovering from unforeseen malfunction(s) or failure.

III. PROPOSED METHOD

In cooperative MARL, both value-based and policy-based techniques are utilized to maximize collective rewards, aiming to converge toward an optimal solution, potentially the global optimum. The convergence process is influenced by the stochastic nature of agents' exploration, particularly in scenarios where multiple cooperation strategies offer the same maximum total reward. Thus, the team's collaboration strategy, and inherently its structural dynamics, are significantly influenced by the randomness in agents' exploration process. However, in real-world applications, the importance of the team's structure becomes pronounced when robots encounter mechanical issues (e.g., battery or joint failures), necessitating a strategy that does not solely depend on randomness. In other words, the current state-of-the-art methods demonstrate limited effectiveness in adapting team behavior to scenarios that involve unforeseen failures.


Algorithm 1: Collaborative Adaptation

input : P-NN, $\hat{Q}^{\text{prediction}}$, T-NN, \hat{Q}^{target} , relational network, G ; batch size, b ; number of iterations for updates, m

- 1 **foreach** *episode* **do**
- 2 Initialize s
- 3 **foreach** *step of episode* **do**
- 4 Choose a from s using policy derived from $\hat{Q}^{\text{prediction}}$ (with ε -greedy)
- 5 Take action a , observe r, s'
- 6 Store s, a, r, s' in memory
- 7 $s \leftarrow s'$
- 8 **for** $i = 1, \dots, m$ **do**
- 9 $S, A, R, S' \leftarrow$ sample chunk, size of b , from memory
- 10 $Q_{\text{values}}^{\text{prediction}} \leftarrow \hat{Q}^{\text{prediction}}(S)$
- 11 $Q_{\text{values}}^{\text{prediction}} \leftarrow$ action A of $Q_{\text{values}}^{\text{prediction}}$ of every agent in every sample
- 12 $Q_{\text{total}}^{\text{prediction}} \leftarrow$ use (10) with G and $Q_{\text{values}}^{\text{prediction}}$
- 13 $Q_{\text{values}}^{\text{target}} \leftarrow \hat{Q}^{\text{target}}(S')$
- 14 $Q_{\text{target}} \leftarrow$ max of $Q_{\text{values}}^{\text{target}}$ of every agent in every sample
- 15 $Q_{\text{total}}^{\text{target}} \leftarrow$ use (10) with G and Q_{target}
- 16 $R^{\text{team}} \leftarrow$ use (9) with G and R
- 17 $L \leftarrow$ use (6) with $R^{\text{team}}, Q_{\text{total}}^{\text{target}}, Q_{\text{total}}^{\text{prediction}}$
- 18 Backpropagate L to the weights of $Q_{\text{total}}^{\text{prediction}}$
- 19 Update the parameters of \hat{Q}^{target} with the parameters of $Q_{\text{total}}^{\text{prediction}}$

To effectively overcome these challenges, we propose a novel framework known as Collaborative Adaptation (CA). This framework is tailored to consider the dynamics between agents, guiding them toward adopting cooperative strategies to improve team performance and accelerate adaptation. In the current field of cooperative MARL, such a mechanism is notably absent, leading to increased difficulty and time required for adaptation in the face of unforeseen malfunctions. Our CA framework aims to fill this gap by encouraging agents, through inter-agent relationships, to either assist malfunctioning one directly or completing its task on behalf of it, thus improving the team's overall performance.

The proposed framework integrates a relational network into the agents' learning process, structured as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$, which defines the inter-agent relationships to enhance cooperative strategies.



Each agent $i \in \{1, 2, \dots, N\}$ is depicted as a vertex v_i in \mathcal{V} , with \mathcal{E} representing the set of directed edges e_{ij} extending from v_i to v_j , and \mathcal{W} containing the edge weights $w_{ij} \in [0, 1]$. These weights, w_{ij} , quantify the level of influence or interest agent i has in the outcomes of agent j .

In the context of the assigned tasks to agents, the relation-

ships may be utilized either for the computation of the team’s action value, Q_{tot} , through the agents’ individual action-values, Q_i , or for the calculation of the team reward, r_{team} , via the individual rewards of the agents, r_i . For instance, in scenarios where agents receive rewards based on individual actions while engaging in a collective task, the aggregate team reward is calculated for use in (6), as follows:

$$r_{\text{team}} = \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{E}_i} w_{ij} r_j. \quad (9)$$

In cases where agents receive a uniform team reward from the environment and lack access to individual rewards, the team’s action-value, Q_{tot} , is re-defined to incorporate the inter-agent relationships for utilization in (6), as:

$$Q_{\text{tot}} = \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{E}_i} w_{ij} Q_j, \quad (10)$$

where Q_j denotes the action-value of agent j .

The modifications introduced in our framework allow for the direction or guidance of the agents’ collaborative approach, thereby improving their adaptability to unexpected robot failure and accelerating the convergence to optimal solutions even without any malfunction, as evidenced by our empirical findings. During our experimental investigations, VDN [20] and Mixed Q-Functionals (MQF) [14] were employed as the learning algorithms for the agents within the CA framework, tailored for discrete and continuous settings, denoted as CA-VDN and CA-MQF, respectively. The pseudo-code corresponding to the CA framework is presented in Algorithm 1.

IV. ENVIRONMENTS

To evaluate how effectively our framework influences agent behaviors and improves their ability to adapt to unexpected failures of their teammate(s) in both discrete and continuous tasks, we conducted experiments in two distinct environments. In the context of discrete tasks, we employed a multi-agent grid-world environment and compared our framework’s performance against value-based methods, particularly IDQN and VDN. For continuous tasks, we used a multi-agent MuJoCo environment and conducted comparisons with policy-based methods, specifically variants of DDPG designed for multi-agent scenarios.

A. Multi-agent Grid-world Environment

In this environment, depicted in Fig 1(a), the goal of each episode is for agents to consume all resources by visiting their locations. Agents can perform five actions: move up, down, left, right, or stay idle. Moreover, they have access to a special action called *push*, which allows them to push adjacent agents. This action is possible when the pushing agent executes a non-idle action towards an idle agent to be pushed. Following a *push*, the agent initiating the push stays stationary, whereas the pushed agent is moved one space in the direction of the push. When an agent successfully consumes a resource, it receives a reward of +10. Each resource is single-use and can be consumed by only one




Fig. 1: (a) multi-agent grid-world environment with four agents. (b-c) Relational networks employed in CA-VDN.

agent. However, the agent incurs an individual penalty of -1 for each time-step per unconsumed resource, except when occupying a resource location, which acts as a safe spot. The episode ends when either all resources have been consumed or the maximum number of time-steps has been achieved.

The design of this environment is intended to be solvable by VDN, while also highlighting the challenges that unforeseen malfunction(s) can introduce, even in simple scenarios. A malfunction is simulated by immobilizing the green agent. This setup illustrates the impact of integrating agent relationships into the learning to address these challenges.

B. Multi-agent MuJoCo

Multi-agent MuJoCo (MAMuJoCo) [27] is a novel benchmark for continuous cooperative multi-agent robotic control. Basically, it is an extension of MuJoCo [28], by creating a wide variety of novel scenarios in which multiple agents within a single robot have to solve a task cooperatively. More specifically, we used the MaMuJoCo-Ant environment which is a robotic simulation environment that emulates a 3D ant, as depicted in Fig.2(b). It consists of a freely rotating torso linked to four legs, with each leg having two joints. The primary goal of this environment is to achieve coordinated movement in the forward direction (toward the positive direction on the x -axis). This is achieved by strategically applying torques to the eight joints, thus controlling the ant’s movements.

In our experiments, we employed a variant of the ant environment that assigns an individual agent to each leg, as




Fig. 2: (a) Representation of an ant featuring four agents, each distinguished by a different color and (b) The MaMuJoCo-Ant simulation environment. (c-d) Relational networks used in CA-MQF.

depicted in Fig2(a). Consequently, the ant is controlled by four distinct agents, each in charge of a leg comprising two joints. In this configuration, each agent has two action values that range from $[-1, +1]$. During an episode, agents receive a team reward at each step, calculated as $r_{\text{team}} = r_{\text{stable}} + r_{\text{forward}} - r_{\text{ctrl_cost}}$, where r_{stable} is a fixed reward given at every time-step for maintaining a stable posture (not upside down), r_{forward} is the reward for moving forward, measured as $\frac{\Delta x}{dt}$ (with dt representing the time between actions and Δx the change in the x direction), and $r_{\text{ctrl_cost}}$ is the penalty for executing excessively large actions. Additionally, agents are penalized with a -100 reward if the ant becomes upside down. For our experiments, we set r_{stable} at $+0.01$ and the maximum number of steps per episode at 100. It is important to notice that agents have no access to their individual rewards, but only team reward, r_{team} .

The integration of multiple agents within a single robotic ant showcases the effectiveness of learning algorithms in enhancing robustness against single-point failures, making it well-suited for our experimental framework. In particular, for our experiments, one leg is immobilized to represent a malfunction. The experiments are designed to evaluate the agents' ability to effectively synchronize actions and adapt to unforeseen malfunction that their teammate(s) might encounter.

V. MODELS AND HYPERPARAMETERS

In this study, the neural network architecture employed for agent modelling in all algorithms tested is the Multi-Layer Perceptron (MLP) for both environments. For the discrete environment, the MLP configuration consists of two hidden layers, each with 128 neurons, using the ReLU activation function. The prediction networks are trained at the end of each episode through 10 iterations, utilizing batches of 32 instances (time-steps) that are randomly selected from the replay memory, capable of storing up to 50k time-steps. Every 200 episodes, the weights of the target neural networks are re-assigned with those of prediction networks. In the continuous environment, the MLPs feature three hidden layers, each with 256 neurons, and employ the TanH activation function (while critic networks in MADDPG still uses the ReLU). The prediction networks' weight are updated every 10 steps with batches of 512, randomly chosen from a replay

memory of 500k time-steps. Target networks receive updates at each time-step through a *soft* update process, integrating a small factor ($\tau = 0.01$) with the weights from the prediction networks. In both discrete and continuous environments, the *Adam* optimizer is utilized for network optimization, and the squared Temporal Difference (TD) error is used as the metric for loss evaluation.

VI. EXPERIMENTAL RESULTS

In this section, we present our experimental results on MaMuJoCo-Ant and our cooperative multi-agent grid-world environment. We have simulated a malfunction after agents converged to one of the several possible solutions (potentially the global optimal solution). Then, we compared the adaptation abilities of our collaborative adaptation framework with IDQN and VDN for discrete tasks, and with IQF and MADDPG for continuous action domains.

In case of a unforeseen failure, our study assumes that a detection mechanism is in place, capable of determining the malfunction's timing (possibly by monitoring the team's collective reward) and identifying the specific agent that is malfunctioning. The identification process might involve analyzing the agents' previous movements to identify deviations from the behavioral patterns established during training.

The CA framework demonstrates superior adaptability to unexpected agent failure compared to IDQN and VDN in our multi-agent grid-world environment involving discrete actions. Despite the initial similarity in performance between the algorithms, IDQN and VDN encounter challenges in recovering post-malfunction, where the green agent, as depicted in Figure 1(a), is immobilized. To improve the algorithms' ability to identify new strategies or behaviors, the epsilon (ϵ) value is reset after the malfunction incident. Fig. 3 displays the experimental outcomes in this setup, showcasing the average training (indicated by shaded areas) and test rewards (represented by solid lines) across ten runs. The test rewards are calculated by assessing the agents' individual rewards using a greedy approach, with the training process being paused after every 50 episodes.

IDQN-trained agents, after malfunction, operate autonomously without a system to foster mutual cooperation, leading to inconsistent behaviors across multiple runs. Fol-




Fig. 3: Multi-agent Grid-world results: Average individual rewards of agents before and after the green agent's malfunction at the 5000th episode for (a) IDQN, (b) VDN and (c) CA-VDN.

lowing the consumption of three resources, agents typically engage in one of three behaviors: (i) the episode continues due to an unconsumed resource, resulting in random actions and penalties for the agents; (ii) an agent may learn to consume the remaining resource, thereby concluding the episode; (iii) an agent might discover how to push the malfunctioning agent, aiming to secure the reward for pushing while another goes to consume the remaining resource to end the episode. However, there is no evidence of intentional collaborative efforts among the agents, such as assisting the malfunctioning agent in consuming the nearest resource by pushing.

In VDN, despite increased exploration following this malfunction, challenges in recovery persist in some runs, leading to unstable agents (refer to Fig. 3(b)). Specifically, in 40% of the runs, agents learn to push malfunctioning agents toward a resource and then proceed to the nearest resource. However, in 60% of the runs, agents either become stuck pushing each other or consume resources without assisting the malfunctioning agent.

On the other hand, the CA framework, employing VDN as its learning algorithm (denoted as CA-VDN), incorporates a relational network as illustrated in Figure 1(c). This network enables all agents to prioritize the green agent, identified as the malfunctioning agent, thereby expediting the adaptation process. In this setting, the relational network is directly integrated into the rewards, given that the agents have access to their individual rewards. The observed collaborative behavior, aimed at optimizing team performance, is characterized by the following movements: (i) the red agent pushes the green agent downward by one cell, (ii) the orange agent moves the green agent one cell right. This cooperation allows the green agent to access the nearest resource location despite its malfunction. The framework’s effectiveness is validated by a comparative analysis of the individual and total rewards of each agent before and after the malfunction, with the results presented in Table I.

The CA framework outperforms both IQF and MADDPG in the MaMuJoCo-Ant environment for continuous action domains. Once the ant is fully capable of navigating the map (up to the 30000th episode), all tested algorithms enable four agents (representing each leg) to collaboratively learn to move the ant in the $+x$ direction. It is crucial to

TABLE I: Average metrics with 95% confidence intervals for ten runs upon training completion.

	Before Malfunction			After Malfunction		
	IDQN	VDN	CA-VDN	IDQN	VDN	CA-VDN
Blue Agent	6.00±0.00	5.80±0.25	5.50±0.64	6.00±2.37	-74.20±18.14	6.90±0.19
Red Agent	6.00±0.00	5.50±0.50	5.70±0.40	-2.40±12.36	-63.60±20.26	9.90±0.19
Orange Agent	6.00±0.00	5.20±0.67	5.40±0.79	8.10±7.39	-66.50±21.47	9.70±0.40
Green Agent	6.00±0.00	5.70±0.40	5.60±0.74	-12.70±6.45	-35.70±28.60	10.50±0.93
Sum	24.00±0.00	22.20±1.64	22.20±2.38	-1.00±22.59	-240.00±72.87	37.00±1.49

note that, although the environment *positively* rewards such movement, it concurrently imposes penalties for the execution of large actions (refer to $r_{\text{ctrl_cost}}$ in IV-B), potentially limiting energy consumption per action.

In this setting, the CA framework employs QF as a learning algorithm, referred to as CA-MQF. Due to the agents’ lack of access to individual rewards, relational networks are applied to their state-action values. As demonstrated in Table II, both CA-MQF, with the relational network depicted in Fig. 2(b), and MADDPG exhibit comparable performance, and they surpass IQF in achieving distance from the origin in the $+x$ direction. However, CA-MQF outperforms both MADDPG and IQF in terms of rewards, as MADDPG tends to generate excessively large actions.

In the 30000th episode, a simulated malfunction restricts the movement of one leg, leading to a significant decrease in the average rewards of the algorithms, as depicted in Fig. 4(a). This incident necessitates modification of the relational network from Fig. 2(b) to Fig. 2(c). The importance placed by non-malfunctioning agents to the malfunctioning one drops to zero (refer to Fig. 2(c)), as its state-action values becomes unreliable because of the malfunction. The trajectories in Fig. 4(b-c) illustrate that CA-MQF, with its adjusted inter-agent relationships, is effective in recovering from such malfunctions. While MADDPG outperforms IQF, it does not reach the level of adaptability demonstrated by CA-MQF. The accompanying video shows the animation of this experiment.¹

Upon deeper analysis, a noticeable difference is observed between the average team reward depicted in Figure 4(a) and Table II. This difference arises because, to reduce the training time, the testing average team reward was computed with only 100 episodes, pausing training every 1000 steps,

¹Accompanying video: <https://youtu.be/-0Qd5jyRGIY>

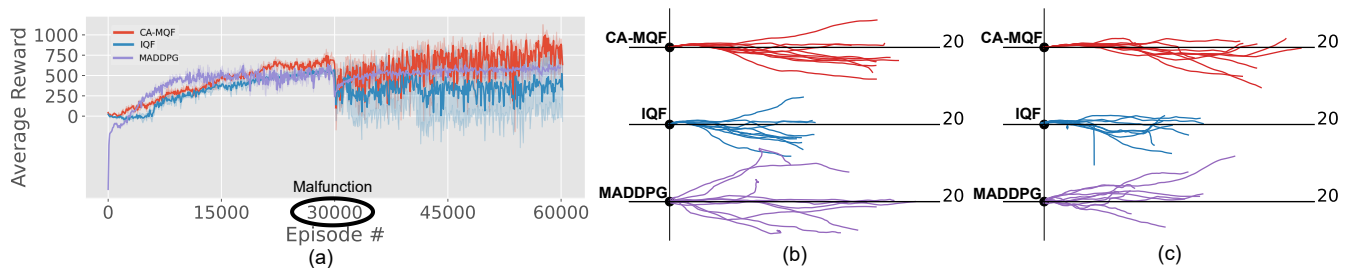


Fig. 4: MaMuJoCo-Ant results: (a) Average team rewards before and after malfunction occurred at the 30000th episode. (b-c) Robot trajectories in x-y plane: (b) before and (c) after malfunction, upon completing 30k and 60k training episodes, respectively. It can be seen that the robot can cover more distance using CA-MQF with higher rewards.

TABLE II: Average metrics with 95% confidence intervals for three runs upon training completion.

	Before Malfunction		After Malfunction	
	Team Reward	Origin Distance	Team Reward	Origin Distance
IQF	568.86 ± 16.06	10.37 ± 0.33	369.95 ± 204.02	8.81 ± 1.65
MADDPG	296.24 ± 15.64	13.81 ± 0.25	340.23 ± 20.02	11.37 ± 0.40
CA-MQF	724.38 ± 19.16	13.20 ± 0.43	768.02 ± 14.55	12.80 ± 0.14

whereas the data in Table II were obtained from 10000 test episodes upon training completion, leading to more reliable results. This may also indicate that the IQF and MADDPG models are not stable, as their rewards decreased when the number of test episodes increased.

Overall, the results from experiments in both the discrete and continuous environments highlight that agents trained using the Collaborative Framework exhibit enhanced cooperative behaviors. Furthermore, these agents display an improved capacity for adapting their policies to recover from unexpected malfunctions, capitalizing on the dynamics of inter-agent relationships.

VII. CONCLUSION AND FUTURE WORK

This study presents a novel framework that integrates inter-agent relationships into the learning process, specifically addressing unexpected malfunction scenarios. The efficacy of our method in improving cooperative behaviors among agents and facilitating effective adaptation to unforeseen robotic malfunctions has been demonstrated in both continuous and discrete environments. The experimental findings indicate that the tested algorithms often stuck at sub-optimal solutions, thus decreasing team rewards, whereas our proposed algorithm successfully mitigates and recovers from these malfunctions. Future research will aim to expand the experiments to more complex situations involving multiple agents encountering various unexpected malfunctions (for instance, noisy or adversarial leg/agent) and undertake a comparative analysis with other leading methodologies to further evaluate the robustness and efficiency of our framework in these settings.

VIII. ACKNOWLEDGMENTS

This work is supported in part by NSF (IIS-2112633) and the Army Research Lab (W911NF20-2-0089).

REFERENCES

- [1] A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey," *Ieee Access*, vol. 6, pp. 28 573–28 593, 2018.
- [2] A. Kleiner, J. Prediger, and B. Nebel, "Rfid technology-based exploration and slam for search and rescue," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, p. 4054.
- [3] S. D. Pendleton, H. Andersen, X. Du, X. Shen, M. Meghiani, Y. H. Eng, D. Rus, and M. H. Ang Jr, "Perception, planning, control, and coordination for autonomous vehicles," *Machines*, vol. 5, no. 1, 2017.
- [4] F. A. Oliehoek, M. T. Spaan, and N. Vlassis, "Optimal and approximate q-value functions for decentralized pomdps," *Journal of Artificial Intelligence Research*, vol. 32, pp. 289–353, 2008.
- [5] M. Hüttenrauch, S. Adrian, G. Neumann, *et al.*, "Deep reinforcement learning for swarm systems," *Journal of Machine Learning Research*, vol. 20, no. 54, pp. 1–31, 2019.
- [6] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.

- [7] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [8] Y. Shoham, R. Powers, and T. Grenager, "If multi-agent learning is the answer, what is the question?" *Artificial intelligence*, vol. 171, no. 7, pp. 365–377, 2007.
- [9] L. Busoniu, R. Babuska, and B. De Schutter, "A comprehensive survey of multiagent reinforcement learning," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 38, no. 2, pp. 156–172, 2008.
- [10] L. Matignon, G. J. Laurent, and N. Le Fort-Piat, "Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems," *The Knowledge Engineering Review*, vol. 27, no. 1, pp. 1–31, 2012.
- [11] R. de Lazaano, K. Andreas, J. J. Tai, S. R. Lee, and J. Terry, "Gymnasium robotics," 2023. [Online]. Available: <http://github.com/Farama-Foundation/Gymnasium-Robotics>
- [12] A. Tampuu, T. Matiisen, D. Kodelja, I. Kuzovkin, K. Korjus, J. Aru, J. Aru, and R. Vicente, "Multiagent cooperation and competition with deep reinforcement learning," *PloS one*, vol. 12, no. 4, 2017.
- [13] P. Sunehag, G. Lever, A. Grusl, W. M. Czarnecki, V. Zambaldi, M. Jaderberg, M. Lanctot, N. Sonnerat, J. Z. Leibo, K. Tuyls, *et al.*, "Value-decomposition networks for cooperative multi-agent learning," *arXiv preprint arXiv:1706.05296*, 2017.
- [14] Y. Findik and S. R. Ahmadzadeh, "Mixed q-functionals: Advancing value-based methods in cooperative marl with continuous action domains," 2024, p. 9.
- [15] R. Lowe, Y. I. Wu, A. Tamar, J. Harb, O. Pieter Abbeel, and I. Mor-datch, "Multi-agent actor-critic for mixed cooperative-competitive environments," *Advances in neural information processing systems*, vol. 30, 2017.
- [16] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, pp. 279–292, 1992.
- [17] S. Lobel, S. Rammohan, B. He, S. Yu, and G. Konidaris, "Q-functionals for value-based continuous control," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 7, 2023.
- [18] M. Tan, "Multi-agent reinforcement learning: Independent vs. cooperative agents," in *Proceedings of the tenth international conference on machine learning*, 1993, pp. 330–337.
- [19] C. Claus and C. Boutilier, "The dynamics of reinforcement learning in cooperative multiagent systems," *AAAI/IAAI*, vol. 1998, no. 746-752, p. 2, 1998.
- [20] T. Rashid, M. Samvelyan, C. S. De Witt, G. Farquhar, J. Foerster, and S. Whiteson, "Monotonic value function factorisation for deep multi-agent reinforcement learning," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 7234–7284, 2020.
- [21] Y. Findik, P. Robinette, K. Jerath, and S. R. Ahmadzadeh, "Impact of relational networks in multi-agent learning: A value-based factorization view," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 4447–4454.
- [22] Y. Findik, H. Osooli, P. Robinette, K. Jerath, and S. R. Ahmadzadeh, "Influence of team interactions on multi-robot cooperation: A relational network perspective," in *2023 International Symposium on Multi-Robot and Multi-Agent Systems*. IEEE, 2023, pp. 50–56.
- [23] S. Lim, A. Joseph, L. Le, Y. Pan, and M. White, "Actor-expert: A framework for using q-learning in continuous action spaces," *arXiv preprint arXiv:1810.09103*, vol. 9, 2018.
- [24] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," *Advances in neural information processing systems*, vol. 12, 1999.
- [25] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," in *International conference on machine learning*. Pmlr, 2014, pp. 387–395.
- [26] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [27] B. Peng, T. Rashid, C. Schroeder de Witt, P.-A. Kamienny, P. Torr, W. Böhm, and S. Whiteson, "Facmac: Factored multi-agent centralised policy gradients," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 208–12 221, 2021.
- [28] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.