
Learning Multiple Coordinated Agents under Directed Acyclic Graph Constraints

Jaeyeon Jang

The Catholic University of Korea

Diego Klabjan

Northwestern University

Han Liu

Northwestern University

Nital S. Patel

Intel, Corporation

Xiuqi Li

Intel, Corporation

Balakrishnan Ananthanarayanan

Intel, Corporation

Husam Dauod

Intel, Corporation

Tzung-Han Juang

Northwestern University

Abstract

This paper proposes a novel multi-agent reinforcement learning (MARL) method to learn multiple coordinated agents under directed acyclic graph (DAG) constraints. Unlike existing MARL approaches, our method explicitly exploits the DAG structure between agents to achieve more effective learning performance. Theoretically, we propose a novel surrogate value function based on a MARL model with synthetic rewards (MARLM-SR) and prove that it serves as a lower bound of the optimal value function. Computationally, we propose a practical training algorithm that exploits new notion of leader agent and reward generator and distributor agent to guide the decomposed follower agents to better explore the parameter space in environments with DAG constraints. Empirically, we exploit four DAG environments including a real-world scheduling for one of Intel’s high volume packaging and test factory to benchmark our methods and show it outperforms the other non-DAG approaches.

1 Introduction

Multi-agent reinforcement learning (MARL) coordinates multiple subtasks to collaboratively achieve an optimal team reward as a shared goal [25]. However, most existing works do not generalize to the settings where multiple subtasks have a complex relationship where higher-level subtasks are affected by lower-level subtasks [23, 6, 15]. In contrast, many real-world systems have nontrivial dependencies among different subtasks. For example, a factory control system must coordinate various processes following precedence constraints [12]. Thus a gap exists between methods and applications. This article aims to propose novel algorithms and theories to bridge this gap.

More specifically, we focus on problems in which subtasks have relationships characterized by a DAG $G := (\mathcal{V}, \mathcal{A})$ where \mathcal{V} and \mathcal{A} denote the set of vertices and the set of arcs, respectively. Arc (u, v) indicates that information flows from u to v such that taking an action for subtask u affects the state of subtask v . We formulate our reinforcement learning (RL) problem as a Markov decision process with DAG constraints (MDP-DAG), defined as the tuple $\mathcal{M} = (\{\mathcal{S}^i | i \in \mathcal{V}\}, \{\mathcal{A}^i | i \in \mathcal{V}\}, \{\mathcal{T}^i | i \in \mathcal{V}\}, \{\mathcal{R}^i | i \in \mathcal{L}\}, \{p_0^i | i \in \mathcal{V}\}, \gamma)$, where \mathcal{L} denotes the set of all sinks in the DAG. Each agent i deals with a subtask in the DAG. The transition dynamic \mathcal{T}^i determines the distribution of the next state s_{t+1}^i given the current state s_t^i and the set of actions $\{a^j | j \in \Delta(i)\}$, where $\Delta(i)$ is the set of nodes in the sub-graph from the source nodes to node i . An agent i for a sink receives a reward $\mathcal{R}^i := r^i(s_t^i, \{a_t^j | j \in \Delta(i)\})$, where $a_t^j \sim \pi^j(\cdot | s_t^j)$ with π^j being the policy for subtask j . Let the initial

state s_0^i be determined by the distribution p_0^i . Then, the objective of learning is to maximize the sum of discounted rewards across all sinks (team rewards), given the structure of the DAG as follows: maximize $\sum_{i \in \mathcal{L}} \mathbb{E}_{\{\pi^j | j \in \Delta(i)\}} [\sum_{t=0}^{\infty} \gamma^t r^i(s_t^i, \{a_t^j | j \in \Delta(i)\})]$, where $\gamma \in [0, 1)$ is the discount factor.

In particular, in the perspective of a high-level subtask, the system does not receive a reward unless all its upstream subtasks have taken actions. Such a delayed rewarding mechanism is common in many real-world problems including industrial process control [9], traffic optimization [8], and resource allocation [22]. Most existing deep reinforcement learning algorithms suffer from inferior performance because no immediate supervision is given in most time steps [7, 10]. In addition, taking into account the complex interactions between different agents, how to distribute the delayed team reward to the agent dealing with different subtasks is quite challenging.

To address these challenges, we first build a theoretical foundation of our approach. Specifically, we prove that we can at least optimize a lower bound of the optimal value function of the DAG system by introducing the concept of synthetic reward. In addition, to ensure practicality, we propose a new training algorithm that introduces two new entities: leader and reward generator and distributor (RGD) as shown in Fig. 1.

In the proposed approach, the leader generates a goal vector for each follower. The goal is not a human-interpretable goal but an abstract signal that evolves during training so that the leader and the followers utilize it together to communicate for a higher achievement. The leader trains the set of goals for better coordination of followers considering the whole environment, and each follower optimizes its policy by pursuing the given goals.

In addition, we introduce the concept of the RGD to coordinate agents in the inner setting, called followers, while considering their contributions to the team rewards in the DAG structure. However, the actual contributions of the agents cannot be easily captured through existing non-DAG MARL approaches. In this paper, we develop a strategy to provide incentives (synthetic rewards) using a RGD that generates and distributes reward so that the followers are guided to explore better. Specifically, if a follower contributes to a high team reward, a high synthetic reward is given to the follower by the RGD. Thus, a follower focuses on optimizing its own policy to obtain a high synthetic reward only based on the state of itself. We believe that the concept of the leader and RGD are introduced the very first time herein to address MDP-DAG.

Our main contributions are as follows.

- We propose MARLM-SR to address MDP-DAG by providing a lower bound of the optimal value function based on team rewards under DAG constraints.
- In the proposed learning algorithm, we introduce a novel leader agent to distribute goals to the followers in the form of simple abstract messages that only the leader and the followers can interpret.
- The concept of the reward generator and distributor is first introduced in the area of reinforcement learning to address the problem of reward shaping in the DAG.
- The proposed learning algorithm demonstrates high practicality and scalability because each follower only needs to consider the state of its own subtask.

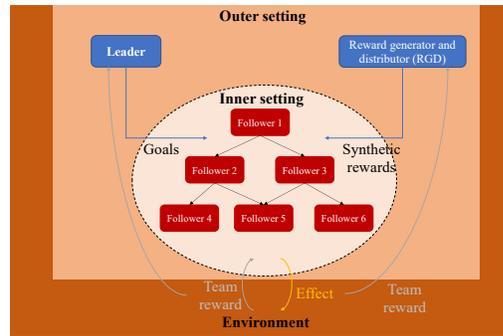


Figure 1: A brief overview of our approach. The leader guides the followers towards higher team rewards by providing agent-specific goals. In addition, the RGD is introduced with the aim of training multiple coordinated followers toward high team rewards. The RGD generates and distributes synthetic rewards so that the followers are rewarded according to their contributions in addition to the team rewards, respectively.

2 Related Works

2.1 Leader-follower MARL

The leader-follower MARL has been proposed to address the problem of coordinating non-cooperative followers by providing goals and/or bonuses to maximize team rewards. However, most existing works focus on simple tabular games or small scale Markov games [20, 16, 14, 2]. Recently, some researchers have proposed deep RL based leader-follower MARL that can be applied to more general problems. For example, Shu and Tian [18] applied deep RL to learn the leader’s policy that assigns goals and bonuses to followers. Similarly, Yu et al. [24] proposed an advanced deep leader-follower MARL algorithm by incorporating a sequential decision module based on the observation that the goal and bonus are sequentially correlated. However, they also have a limitation of assuming homogeneous followers. More seriously, to the best of our knowledge, there is no leader-follower approach or MARL algorithm that can be used to coordinate multiple agents in a DAG, which is our target.

2.2 Reward shaping for multi-agent systems

Often, environmental feedback is not enough to effectively train an agent, especially when the environment is stochastic [4]. In this case, reward shaping can help guide an agent’s exploration by providing an additional artificial reward signal. A few researchers have proposed reward shaping methods for multi-agent systems. Colby et al. [3] showed that their algorithm called ‘difference rewards’ is powerful in effectively allocating rewards across multiple agents. Here, ‘difference rewards’ was designed to reveal the contribution of the current action of an agent by comparing the current reward to the reward received when the action of the agent is replaced with the default action [21]. In practice, ‘difference rewards’ can be estimated using a function approximation technique [6]. It has been proven that potential-based reward shaping, which is one of the typical reward shaping methods, does not alter the optimal policy [13]. Based on this background, Devlin et al. [5] proposed two potential-based reward shaping methods based on ‘difference rewards.’ Even though this algorithm guarantees optimality, it assumes top-down MARL, in which all agents have a common task and a centralized system distributes rewards to the agents based on their contributions. Thus, it lacks scalability and applicability. To tackle this problem, Aotani et al. [1] proposed a localized reward shaping method that prevents the agents from knowing the interests between them. However, this work still cannot consider the relationship between agents in a DAG.

3 Modeling Setting

Global decision-making is mainly used for many real-world systems. However, traditional global single-agent RL models (GSARLMs) are poorly suited to environments under DAG constraints even though the global model can provide an optimal or a very good solution theoretically [11]. This is because, in general, the search space for obtaining a single global solution is too large while compromising scalability. In addition, GSARLM cannot easily capture interactions between multiple subtasks in a DAG. Thus, in this section, we define the MARL model with synthetic rewards (MARLM-SR) and build an analytical background. In addition, we further decompose the problem by introducing the concept of goal periods. Finally, we provide strong evidence of higher practicality and scalability of MARLM-SR based on this decomposed problem by proposing a training algorithm in the next section.

3.1 MARLM-SR

The objective of GSARLM is to derive an optimal solution that covers all subtasks considering the current states of all subtasks altogether. Even though one action is made to cover all subtasks, the state transition of each subtask is stochastically determined based on inherent DAG relationships. Let s_t^i and a_t^i be the state and action of subtask $i \in \mathcal{V}$. First, the lowest-level subtasks, the source nodes i in the DAG, are affected only by themselves based on the stochastic state transition $s_{t+1}^i \sim p(\cdot | s_t^i, a_t^i)$. On the other hand, the states of the other subtasks are affected by the ancestor nodes in the DAG, $s_{t+1}^j \sim p(\cdot | s_t^i, \{a_t^j | j \in \Delta(i)\})$.

Let us assume that Π , the policy for the entire system, can be decomposed into $(\pi^1, \pi^2, \dots, \pi^I)$ in which π^i is the policy for subtask i , where $I = |\mathcal{V}|$. In addition, since the performance of a system

with a DAG structure is represented by the rewards of the sinks in the DAG, the highest-level subtasks, we assume that the team reward is the sum of the rewards obtained from sinks. Let r^i be the reward function of subtask i , $i \in \mathcal{L}$, the set of all sinks. Then, we define the value function of a subtask i in \mathcal{L} as follows

$$V_i^{\{\pi^j | j \in \Delta(i)\}}(s_0^i) = \mathbb{E}_{\{\pi^j | j \in \Delta(i)\}} \left[\sum_{t=0}^{\infty} \gamma^t r^i(s_t^i, \{a_t^j | j \in \Delta(i)\}) \right] \quad (1)$$

where $V_i^{\{\pi^j | j \in \Delta(i)\}}$, the value of subtask i , has dependency on $\{\pi^j | j \in \Delta(i)\}$. The objective function is maximize $\sum_{i \in \mathcal{L}} V_i^{\{\pi^j | j \in \Delta(i)\}}(s_0^i)$.

Next, we introduce the concept of MARL with synthetic rewards. First, an agent deals with its own subtask and receives a synthetic reward. Here, we assume that the synthetic reward for an agent is determined by considering its contribution to the team rewards. In other words, an agent's policy which contributes to a high reward of its descendant sinks yields a high synthetic reward. We assume that there can be a function f_{ik} that measures the contribution of agent i to sink agent k 's reward and the total contribution of agents in $\Delta(k)$ to sink agent k 's reward is less than or equal to 1 as shown in (2) because the reward of a sink agent is also affected by environmental feedbacks. All subtasks that have a path to/from subtask i have an impact on the agent i 's contribution. Thus, the synthetic reward function of agent i has dependency on $\Omega(i) = \Delta(i) \cup \Upsilon(i)$, where $\Upsilon(i)$ denotes the set of subtasks in the induced sub-graph rooted in subtask i including node i . Finally, we have the following definition.

Definition 1 Let f_{ik} be a function that produces the magnitude of agent i 's contribution to sink agent k 's reward for $k \in \Upsilon(i)$. For any f_{ik} satisfying

$$\sum_{i \in \Delta(k)} f_{ik}((s_t^j, a_t^j) | j \in \Delta(k)) \leq 1 \quad \forall k \in \mathcal{L}, \quad (2)$$

the synthetic reward function sr^i of subtask i is defined as

$$sr^i((s_t^j, a_t^j) | j \in \Omega(i)) = \sum_{k \in \mathcal{L} \cap \Upsilon(i)} f_{ik}((s_t^j, a_t^j) | j \in \Delta(k)) r^k(s_t^k, \{a_t^j | j \in \Delta(k)\}) \quad \forall i \in \mathcal{V}. \quad (3)$$

Definition 2 We define synthetic value functions based on synthetic rewards as

$$\tilde{V}_i^{\{\pi^j | j \in \Omega(i)\}}(s_0^i) = \mathbb{E}_{\{\pi^j | j \in \Omega(i)\}} \left[\sum_{t=0}^{\infty} \gamma^t sr^i((s_t^j, a_t^j) | j \in \Omega(i)) \right] \quad \forall i \in \mathcal{V}. \quad (4)$$

Next, we show that the total synthetic value provides a lower bound on the total value; thus, we can optimize agents' policies such that synthetic values are maximized in order to maximize a lower bound of the sum of optimal values. It provides the theoretical background that we only need to train agents to seek high synthetic rewards in a parallel fashion. In Section 4, we propose a practical algorithm for generating and distributing synthetic rewards.

Theorem 1 If reward $r^i \geq 0$, $\forall i \in \mathcal{L}$, then, for any f_{ik} satisfying (2), we have

$$\sum_{i \in \mathcal{V}} \tilde{V}_i^{\{\pi^j | j \in \Omega(i)\}}(s_0^i) \leq \sum_{i \in \mathcal{L}} V_i^{\{\pi^j | j \in \Delta(i)\}}(s_0^i). \quad (5)$$

Proof. A detailed proof of this theorem is given in the supplementary material.

3.2 MARLM-SR with goal period

We further extend MARLM-SR by introducing the notion of a goal period, which is a short interval that partitions an episode, enabling more refined coordination between agents over the learning process using two novel entities: leader and RGD. Let D be the number of steps for a goal period, and s_{ld}^i and a_{ld}^i be the state and action at d -th step in l -th goal period, respectively. As a consequence (1) and (4) change to

$$V_i^{\{\pi^j | j \in \Delta(i)\}}(s_{01}^i) = \mathbb{E}_{\{\pi^j | j \in \Delta(i)\}} \left[\sum_{l=0}^{\infty} \sum_{d=1}^D \gamma^{lD+d-1} r^i(s_{ld}^i, \{a_{ld}^j | j \in \Delta(i)\}) \right] \quad \forall i \in \mathcal{L} \quad (6)$$

and

$$\tilde{V}_i^{\{\pi^j|j \in \Omega(i)\}}(s_{01}^i) = \mathbb{E}_{\{\pi^j|j \in \Omega(i)\}} \left[\sum_{l=0}^{\infty} \sum_{d=1}^D \gamma^{lD+d-1} sr^i((s_{ld}^j, a_{ld}^j)|j \in \Omega(i)) \right] \forall i \in \mathcal{V}, \quad (7)$$

respectively. From Theorem 1, (6), and (7), we obtain

$$\max_{\{f_{ik}|k \in \mathcal{L}, i \in \Delta(k)\}} \sum_{i \in \mathcal{V}} \tilde{V}_i^{\{\pi^j|j \in \Omega(i)\}}(s_{01}^i) \leq \sum_{i \in \mathcal{L}} V_i^{\{\pi^j|j \in \Delta(i)\}}(s_{01}^i), \quad (8)$$

subject to $\{f_{ik}|k \in \mathcal{L}, i \in \Delta(k)\}$ complying to Definition 1. This is the basis of our algorithm presented in the next section.

4 Algorithm

In this section, we describe the training algorithm for MARLM-SR. The algorithm consists of the outer and inner settings. In the inner setting, the followers perform their subtasks given by the defined DAG every time step. On the other hand, in the outer setting, two different types of agents are trained to guide the followers to achieve a high team reward. If the followers are guided well based on the policies of the outer agents and a high team reward is achieved, this high team reward is given to the outer agents. We provide a more detailed exposition of the algorithm in the supplementary material.

4.1 Outer setting

The leader provides a different goal to each follower at the beginning of each goal period. It is governed by an RL model with policy π^L . Here, the goal is a vector with fixed length in which each element has a value between 0 and 1. It is used for communication between the leader and the followers. Since the leader is rewarded based on the followers' achievements, it must be trained to produce meaningful goals. On the other hand, the followers must interpret the goals and use this information to achieve high team rewards. Let $S_{ld} = (s_{ld}^i|i \in \mathcal{V})$ be the global state at step d and $G_l = (g_l^i|i \in \mathcal{V})$ be the set of goals in the l -th goal period. Each follower augments its state with g_l^i and thus the state of follower i at step d is $\bar{s}_{ld}^i = (s_{ld}^i, g_l^i)$. In addition, the RGD is modeled with policy π^{RGD} that produces synthetic reward sr^i for each follower i after the l -th goal period (details for generating sr^i are provided later in this section).

The leader is trained to produce G_l that maximizes team rewards since the team rewards are also given to the leader as its own reward. The leader receives cumulative team rewards after each goal period. Thus, the reward of the leader after the l -th goal period is defined as $\sum_{i \in \mathcal{L}} \sum_{d=1}^D r^i(s_{ld}^i, \{a_{ld}^j|j \in \Delta(i)\})$. By extending this cumulative reward to cover infinite goal periods, the objective function for the leader is defined as

$$\begin{aligned} & \underset{\pi^L}{\text{maximize}} V_L^{\{\pi^L, \pi^{RGD}, \pi^j|j \in \Delta(i)\}}(S_{01}) = \\ & \sum_{i \in \mathcal{L}} \mathbb{E}_{\{\pi^L, \pi^{RGD}, \pi^j|j \in \Delta(i)\}} \left[\sum_{l=0}^{\infty} \gamma^l \sum_{d=1}^D r^i(s_{ld}^i, \{a_{ld}^j|j \in \Delta(i)\}) \right], \end{aligned} \quad (9)$$

where the state transition of S_{ld} (in a particular goal period) depends on the underlying policies. The state of the leader is defined as $S_l^L = S_{l1} \circ (g_{l-1}^i|i \in \mathcal{V}) \circ (sr_{l-1}^i|i \in \mathcal{V})$, including the initial global state in each goal period l . By \circ we denote the concatenation operator. Then, the state transition of the leader is defined as $S_{l+1}^L \sim p(\cdot|S_l^L, \{a_{ld}^i|i \in \mathcal{V} \text{ and } d = 1, \dots, D\}, \{g_l^i|i \in \mathcal{V}\}, \{sr_l^i|i \in \mathcal{V}\}) \circ (g_l^i|i \in \mathcal{V}) \circ (sr_l^i|i \in \mathcal{V})$. Additionally, the set of goals are produced based on $(g_l^i|i \in \mathcal{V}) \sim \pi^L(\cdot|S_l^L)$.

The RGD should be able to figure out the followers' state changes to provide effective coordination strategies. The easiest way is to collect the global state for all time steps in a goal period and use it as the input state. However, to prevent the RGD's input from being too high dimensional, we sample global states with equal time step intervals including the first and last global states in a goal period. For simplicity, we call the set of sampled global states as the global state flow (GSF). This state GSF is defined as $gsf_l = (S_{l, kj+1}|j = 0, \dots, \lfloor \frac{D-1}{k} \rfloor) \circ S_{l+1, 1}$, where k is a hyperparameter and $\lfloor \cdot \rfloor$ is the floor function. Vector $S_{l+1, 1}$ is the global state after the last action set $\{a_{l, D}^i|i \in \mathcal{V}\}$ is taken in the l -th goal period. Goals are also used to guide the RGD; thus, the state of the RGD is

$S_l^{RGD} = gsf_l \circ (g_l^i | i \in \mathcal{V})$. The state transition of the RGD is defined as $S_{l+1}^{RGD} \sim p(\cdot | gsf_l, \{a_{l+1,d}^i | i \in \mathcal{V} \text{ and } d = 1, \dots, D\}, \{g_{l+1}^i | i \in \mathcal{V}\}, \{sr_l^i | i \in \mathcal{V}\}) \circ \{g_{l+1}^i | i \in \mathcal{V}\}$.

The RGD policy produces a team reward signal q_l , node values $(v_l^i | i \in \mathcal{V})$, and arc values $(e_l^{(i,j)} | (i,j) \in \mathcal{A})$ for synthetic reward generation and distribution. All these values are within the range $[0, 1]$. The policy is specified by

$$(q_l) \circ (v_l^i | i \in \mathcal{V}) \circ (e_l^{(i,j)} | (i,j) \in \mathcal{A}) \sim \pi^{RGD}(\cdot | S_l^{RGD}). \quad (10)$$

Vector $(sr_l^i | \forall i \in \mathcal{V})$ is obtained based on $(q_l) \circ (v_l^i | i \in \mathcal{V}) \circ (e_l^{(i,j)} | (i,j) \in \mathcal{A})$, not by a closed-form function, but by the proposed reward generation and distribution algorithm exhibited next.

The synthetic reward sr_l^i , $i \in \mathcal{V}$, is given to the followers as a bonus after each goal period. The RGD should provide a high synthetic reward if followers use policies that lead to high team rewards. In addition, the value of the synthetic reward must be adjusted dynamically to make the policy of the RGD significant. This is because followers are more likely to achieve higher team rewards as training progresses. In this case, the same reward can be too small for followers who have had enough training but can be too large for followers without enough training. The quality of the learned policy is revealed as the team reward of the previous episode. The RGD policy produces q_l (in addition to v and e). This value is multiplied with $\frac{\bar{R}_e}{\bar{N}_e}$, the average team reward per goal period, in the previous episodes, where \bar{N}_e is the average number of goal periods and \bar{R}_e is the average total team reward. Finally, in the current episode, the total synthetic reward after the l -th goal period is $M_l = q_l \frac{\bar{R}_e}{\bar{N}_e}$. We simply set $\bar{R}_0 = 0$ or a negligible value.

We assume that the synthetic reward for the follower i is determined based on its contributions to the sink followers among its descendants and their rewards as defined in (3). Thus, we propose a synthetic reward distribution strategy that first sets synthetic reward portions for the followers in sinks considering their achievements, and then sends them down to account for the contributions of lower-level followers. The RGD is trained to achieve high team rewards by creating a good distribution strategy because it is quite challenging to estimate the exact contribution of each agent.

The RGD distributes the synthetic reward generated by the reward generator as shown in Fig. 2. Because the synthetic reward flows in the opposite direction of the task flow, arc (i, j) denotes a directed edge from a higher-level node i to a lower-level node j . We can sequentially calculate shares from the highest-level to the lowest-level followers. First, we calculate initial share \tilde{sh}_l^i for a highest-level follower $i \in \mathcal{L}$ after goal period l as follows

$$\tilde{sh}_l^i = \begin{cases} \frac{v_l^i}{\sum_{k \in \mathcal{L}} v_l^k}, & \text{if } \sum_{k \in \mathcal{L}} v_l^k > 0 \\ \frac{1}{|\mathcal{L}|}, & \text{otherwise} \end{cases} \quad (11)$$

where $|\mathcal{L}|$ is the number of the sinks. Similarly, for each follower, the initial share can be determined after receiving all the rewards from one level higher followers. After all children of the agent i determine the share to the agent i , the initial share \tilde{sh}_l^i is simply calculated by $\tilde{sh}_l^i = \sum_{k \in ch(i)} sh_l^{k,i}$, where $sh_l^{k,i}$ is the share from k to i . After \tilde{sh}_l^i is determined, the final reward shares to the follower i itself and the arc (i, j) are defined as (12) and (13), respectively. Here, $\delta(i)$ denotes the parents of the follower i .

$$sh_l^i = \begin{cases} \tilde{sh}_l^i \times \frac{v_l^i}{v_l^i + \sum_{j \in \delta(i)} e_l^{i,j}}, & \text{if } v_l^i + \sum_{j \in \delta(i)} e_l^{i,j} > 0 \\ \frac{1}{1+|\delta(i)|}, & \text{otherwise} \end{cases} \quad (12)$$

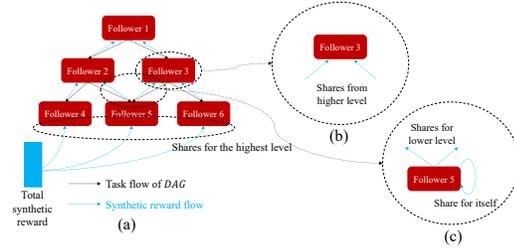


Figure 2: An overview of reward distribution. The algorithm first determines the shares for the highest-level followers. Then, one level lower followers receive the shares set as shown in (b). After receiving the rewards in (c), the shares for itself and for the one level lower followers are determined. This process is repeated until the lowest-level followers, root nodes, receive their shares.

$$sh_l^{i,j} = \begin{cases} \tilde{sh}_l^i \times \frac{e_l^{i,j}}{v_l^i + \sum_{j \in \delta(i)} e_l^{i,j}}, & \text{if } v_l^i + \sum_{j \in \delta(i)} e_l^{i,j} > 0 \\ \frac{1}{1+|\delta(i)|}, & \text{otherwise} \end{cases} \quad (13)$$

After sh_l^i is determined for all $i \in \mathcal{V}$, $sr_l^i = sh_l^i M_l$ is provided to agent i as the synthetic reward after the goal period l .

Same as the leader, the RGD is trained with the aim of maximizing team rewards by obtaining better coordination through synthetic rewards. However, since the first action of the RGD is taken after the first goal period, we define the value function for the RGD as

$$V_{RGD}^{\{\pi^L, \pi^{RGD}, \pi^j | j \in \Delta(i)\}}(gsf_0) = \sum_{i \in \mathcal{L}} \mathbb{E}_{\{\pi^L, \pi^{RGD}, \pi^j | j \in \Delta(i)\}} \left[\sum_{l=1}^{\infty} \gamma^{l-1} \sum_{d=1}^D r^i(s_{ld}^i, \{a_{ld}^j | j \in \Delta(i)\}) \right], \quad (14)$$

and train the RGD to maximize it.

4.2 Inner setting

In the inner setting, the followers are trained with the supervision of the outer agents. Because the goal given by the leader is incorporated into the state, state transition is defined as $\bar{s}_{l,d+1}^i \sim p(\cdot | \bar{s}_{ld}^i, \{a_{ld}^j | j \in \Delta(i)\})$. In each episode during training, followers' achievements are rewarded in two ways. First, the followers share the team reward equally because it is not only quite challenging to create synthetic rewards based on the exact contribution to the team reward, but the team reward can also serve as effective supervision. For each follower, $\frac{\sum_{i \in \mathcal{L}} r^i(s_{ld}^i, \{a_{ld}^j | j \in \Delta(i)\})}{|\mathcal{V}|}$ is given as a shared team reward at the d -step of the l -th goal period. In addition, the follower i receives a synthetic reward sr_l^i from the RGD after the l -th goal period based on the difference in their achievements. By considering both the shared team reward and the synthetic reward, we define the objective function of the follower i as

$$\begin{aligned} & \underset{\pi^i}{\text{maximize}} \bar{V}_i^{\{\pi^L, \pi^{RGD}, \pi^j | j \in \mathcal{V}\}}(\bar{s}_{01}^i) = \\ & \mathbb{E}_{\{\pi^L, \pi^{RGD}, \pi^j | j \in \mathcal{V}\}} \left[\sum_{l=0}^{\infty} \left[\gamma^{(l+1)D-1} sr_l^i + \sum_{d=1}^D \sum_{k \in \mathcal{L}} \gamma^{lD+d-1} \frac{r^k(s_{ld}^k, \{a_{ld}^u | u \in \Delta(k)\})}{|\mathcal{V}|} \right] \right]. \end{aligned} \quad (15)$$

Here, we use \bar{V} to distinguish it from the value functions in the modeling section, which only consider explicit rewards or synthetic rewards.

In the algorithm, the leader sets goals at the beginning of each goal period and is rewarded after the goal period. On the other hand, the RGD determines the synthetic reward distribution strategy after each goal period. And this strategy influences the followers to behave differently in the next goal periods. Therefore, the RGD is rewarded in the next goal period.

5 Experiments

5.1 Implementation details

We used a proximal policy optimization algorithm [17] to optimize the policies of all agents used in this work. The hyperparameters used for the proposed algorithm and the other baselines are summarized in the supplementary material. In our implementation, we included only the initial global state S_{l1} at each goal period l as the input state vector for the leader, to enhance tractability. For RGD, we introduced two separate networks using the same state vector; one for the reward generator that produces q_l , and another for the reward distributor that creates $(v_l^i | i \in \mathcal{V})$, $(e_l^{i,j} | (i,j) \in \mathcal{A})$, and finally generates synthetic rewards. To obtain synthetic rewards, we first need to calculate the average team reward per goal period, $\frac{\bar{R}_e}{N_e}$. In experiments, we only considered the number of goal periods and the total team reward from the immediately preceding episode. All algorithms compared in this work were implemented based on the TensorFlow framework.

5.2 Environments

We created three artificial environments to simulate systems with DAG constraints: a factory production planning case, a logistics case, and a hierarchical predator-prey case. We also investigated the performance of the proposed algorithm in real-world scheduling for one of Intel’s high volume packaging and test factories. The details of all environments are described in the supplementary material.

5.3 Baselines

We have compared seven baseline algorithms against our algorithm. First, we used the following five algorithms that do not employ reward shaping.

- Global single-agent algorithm (GS): In this baseline, a single agent is learned to do all subtasks.
- Shared reward multi-agent algorithm (SRM): Each agent deals with a subtask and shares the reward. This algorithm is perhaps the most popular multi-agent learning algorithm, also known as independent Q-learning [19] or independent actor-critic [6], depending on the type of the learner used.
- Leader-follower multi-agent algorithm (LFM): This baseline adds the leader to SRM. Specifically, the followers are given the goals as well as the shared rewards.
- RGD-follower multi-agent algorithm (RFM): The RGD is added to SRM in this baseline. Thus, the followers are given the synthetic rewards as well as the shared team rewards.
- The proposed algorithm: We have the leader, the RGD, and the followers of the proposed algorithm. This algorithm adds the RGD to LFM and the leader to RFM.

The last two are our stripped-down algorithms and, as such, not previously existing algorithms. We are not aware of any reward shaping method targeting DAGs, however we found two existing reward shaping methods that can be applied to coordinate multiple agents. We introduced these two reward shaping methods to the MARL algorithm that trains agents in parallel. Specifically, we also compared the following two baselines against ours.

- Difference rewarding method [3] + MARL algorithm (Diff-M)
- Counterfactual as Potential [5] + MARL algorithm (CaP-M)

5.4 Results

In our proposed algorithm, the outer agents are trained to coordinate followers by providing additional synthetic rewards that correspond to the contributions of the followers in the DAG. To ascertain the effectiveness of reward shaping, we initially evaluated the proposed algorithm against Diff-M and CaP-M. Fig. 3 shows comparison results on the three artificial benchmark cases. The plots use the moving window method, which averages the team rewards over 100 episodes with a step size of one, to reduce variability. The results demonstrate that our method achieves significantly superior performance across all three benchmark cases. Specifically, in terms of the average team reward over the last 100 episodes for the three artificial cases, the proposed algorithm achieves performance that is 132.7% and 89.3% higher than that of Diff-M and CaP-M, respectively. This suggests that, until now, there has not been an effective reward shaping method for systems under DAG constraints. In the case of logistics, ours quickly get away from a bad local optima where agents send almost nothing to the next level agents to reduce inventory cost (refer to the description about the logistics case in the supplementary material), even after it get stuck in.

We also compared the two baseline algorithms across diverse scheduling scenarios. Specifically, we trained the agents in the DAG using the proposed algorithm, Diff-M, and CaP-M and then evaluated their performance on 1,000 new scheduling scenarios (episodes). Fig. 4 presents the histogram comparing the completion rates of the three baselines. In the histograms, we omitted the labeling of x-axis values for confidentiality reasons; however, all histograms share the same scale, with equally spaced intervals along the x-axis. From the results it is clear that our proposed algorithm achieves higher overall completion rates. Specifically, the proposed algorithm demonstrated a performance improvement of 19.2% and 4.4% in terms of the mean completion rate, compared to Diff-M and CaP-M, respectively. In summary, our proposed method of synthetic reward generation and distribution, coupled with communication through the leader’s goals, can enhance coordination leading to increased team rewards.

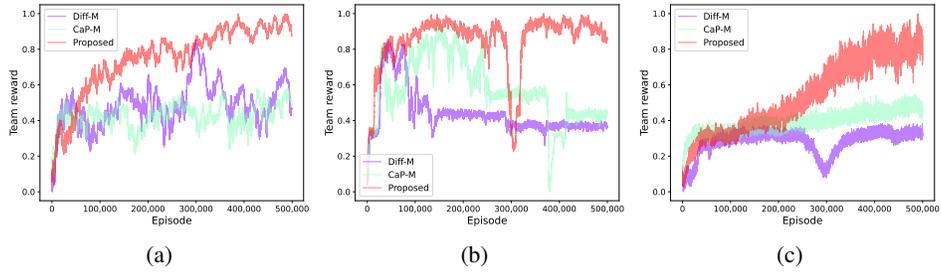


Figure 3: Comparison with state-of-the-art algorithms on (a) the factory production planning case, (b) the logistics case, and (c) the hierarchical predator-prey case. Min-max normalization is applied to the team reward to standardize the scale of the y-axis across the three cases.

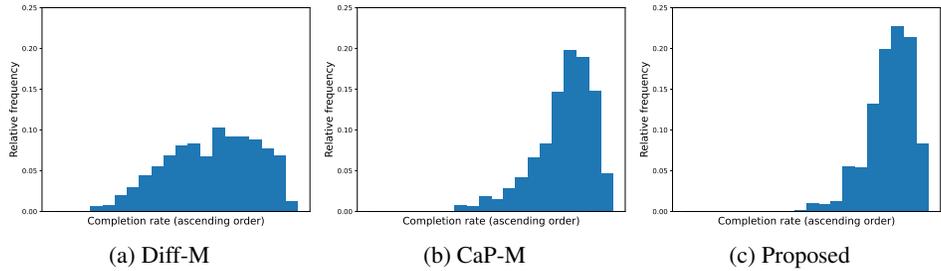


Figure 4: The histogram of the completion rate over 1,000 scheduling scenarios (episodes) for comparison with the state-of-the-art algorithms.

We also conducted ablation studies to evaluate the effectiveness of each component in the proposed algorithm. Fig. 5 shows the comparison results of the five baselines: GS, SRM, LFM, RFM, and our proposed algorithm, on the three artificial benchmark cases. Specifically, GS shows the worst performance in all three cases, revealing that introducing the multi-agent concept is effective for environments with DAG constraints. The leader can help improve performance as shown in (b) and (c). However, by comparing LFM and RFM, we find that the RGD contributes more to performance improvement than the leader in (a) and (c). Specifically, on average over the three cases, LFM and RFM improve the average team reward over the last 100 episodes by 5.6% and 56.0% compared to SRM, respectively. Nonetheless, the proposed algorithm demonstrates the best learning curve in all settings, while achieving an 82.4% higher average team reward compared to SRM. In addition, the performances of LFM and RFM in Fig. 5 are overall better than those of Diff-M and CaP-M in Fig. 3. In other words, we are able to achieve better performance only by adding one component, either the leader or the RGD, in DAG environments. In addition, combining the two components further enhances performance.

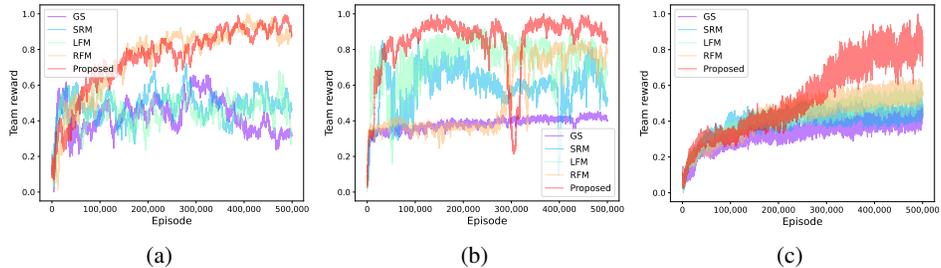


Figure 5: Learning curves of five baselines for ablation study on (a) the factory production planning case, (b) the logistics case, and (c) the hierarchical predator-prey case. Min-max normalization is applied to the team reward to standardize the scale of the y-axis across the three cases.

We also compared the five baselines in diverse scheduling scenarios. The histogram of the completion rate for the five baselines can be found in the supplementary material. We conducted statistical significance tests to validate whether the proposed algorithm is significantly better than the other baselines in terms of the completion rate as shown in Table 1. In the table, we report the average improvement of the completion rate compared to GS instead of the mean completion rate values for confidentiality. The result demonstrates the significant superiority of the proposed algorithm over the other baselines except for LFM. The proposed algorithm achieves a performance improvement of 3.9% by introducing the RGD, and an improvement of 8.5% by introducing both the leader and the RGD together. Even though LFM achieved a good performance similar to ours, the contribution of the RGD is not negligible considering the results in Fig. 5. Thus, we can state both the leader and the RGD are necessary for our algorithm.

Table 1: Result of statistical significance tests.

	GS	SRM	LFM	RFM	Proposed
Mean improvement	-	0.464	0.529	0.494	0.529
p-value	<0.001	<0.001	0.555	<0.001	-

6 Discussion

In this paper, a theoretical background on MARLM-SR was established and a novel training algorithm for coordinating multiple agents in a DAG environment was proposed. Comparison results in several DAG environments including a real-world scheduling environment confirmed that our approach significantly outperforms existing algorithms for DAG systems. It was found that the leader and the RGD contributed to this overwhelming performance.

6.1 Limitations

One limitation of this work is that we did not provide a mathematical basis for whether the synthetic reward obtained through our algorithm satisfies the conditions in the modeling section. Instead, the superiority of the proposed algorithm was shown through empirical results. Nonetheless, there have been few opportunities to apply our algorithm to real-world industrial cases. Therefore, in future studies, the proposed algorithm will be further developed by applying it to more diverse real-world industrial cases.

6.2 Broader impacts

This work seeks to contribute to the MARL modeling and training algorithm for DAG systems. The social impact of the proposed model and algorithm is difficult to predict because it can be utilized universally. However, the results of this study have a positive impact on a variety of industrial applications.

References

- [1] T. Aotani, T. Kobayashi, and K. Sugimoto. Bottom-up multi-agent reinforcement learning by reward shaping for cooperative-competitive tasks. *Applied Intelligence*, 51(7):4434–4452, 2021.
- [2] C. Cheng, Z. Zhu, B. Xin, and C. Chen. A multi-agent reinforcement learning algorithm based on Stackelberg game. In *IEEE Data Driven Control and Learning Systems Conference*, pages 727–732, 2017.
- [3] M. Colby, W. Curran, and K. Tumer. Approximating difference evaluations with local information. In *International Conference on Autonomous Agents and Multiagent*, 2015.
- [4] S. Devlin and D. Kudenko. Plan-based reward shaping for multi-agent reinforcement learning. *The Knowledge Engineering Review*, 31(1):44–58, 2016.
- [5] S. Devlin, L. Yliniemi, D. Kudenko, and K. Turner. Potential-based difference rewards for multiagent reinforcement learning. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 165–172, 2014.
- [6] J. Foerster, G. Farquhar, T. Afouras, N. Nardelli, and S. Whiteson. Counterfactual multi-agent policy gradients. In *AAAI Conference on Artificial Intelligence*, 2018.

- [7] T. Gangwani, Q. Liu, and J. Peng. Learning self-imitating diverse policies. *International Conference on Learning Representations (ICLR)*, pages 1–18, 2019.
- [8] Y. Gong, M. Abdel-Aty, Q. Cai, and M. S. Rahman. Decentralized network level adaptive signal control by multi-agent deep reinforcement learning. *Transportation Research Interdisciplinary Perspectives*, 1:100020, 2019.
- [9] D. Hein, S. Depeweg, M. Tokic, S. Udluft, A. Hentschel, T. A. Runkler, and V. Sterzing. A benchmark environment motivated by industrial control problems. In *IEEE Symposium Series on Computational Intelligence*, pages 1–8, 2018.
- [10] Y. Liu, Y. Luo, Y. Zhong, X. Chen, Q. Liu, and J. Peng. Sequence modeling of temporal credit assignment for episodic reinforcement learning. *arXiv:1905.13420*, 2019.
- [11] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pages 6382–6393, 2017.
- [12] M. Milosevic, R. Cep, L. Cepova, D. Lukic, A. Antic, and M. Djurdjev. A hybrid grey wolf optimizer for process planning optimization with precedence constraints. *Materials*, 14(23):7360, 2021.
- [13] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations : Theory and application to reward shaping. In *International Conference on Machine Learning*, volume 3, pages 278–287, 1999.
- [14] R. Sabbadin and A. F. Viet. Leader-Follower MDP models with factored state space and many followers - Followers abstraction, structured dynamics and state aggregation. In *European Conference on Artificial Intelligence*, pages 116–124, 2016.
- [15] T. Rashid, M. Samvelyan, C. Schroeder, G. Farquhar, J. Foerster, and S. Whiteson. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 4295–4304, 2018.
- [16] R. Sabbadin and A. F. Viet. A tractable Leader-Follower MDP model for animal disease management. In *AAAI Conference on Artificial Intelligence*, pages 1320–1326, 2013.
- [17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [18] T. Shu and Y. Tian. M³RL: Mind-aware multi-agent management reinforcement learning. In *International Conference on Learning Representations*, 2019.
- [19] M. Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *International Conference on Machine Learning*, pages 330–337, 1993.
- [20] K. Tharakunnel and S. Bhattacharyya. Leader-follower semi-Markov decision problems: Theoretical framework and approximate solution. In *IEEE Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pages 111–118, 2007.
- [21] D. H. Wolpert and K. Tumer. Optimal payoff functions for members of collectives. *Advances in Complex Systems*, 4:265–279, 2001.
- [22] Z. Xu, Z. Li, Q. Guan, D. Zhang, Q. Li, J. Nan, C. Liu, W. Bian, and J. Ye. Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 905–913, 2018.
- [23] J. Yang, A. Nakhaei, D. Isele, K. Fujimura, and H. Zha. CM3: Cooperative multi-goal multi-stage multi-agent reinforcement learning. In *International Conference on Learning Representations*, 2020.
- [24] R. Yu, X. Wang, R. Wang, Y. Zhang, B. An, Z. Y. Shi, and H. Lai. Learning expensive coordination: An event-based deep RL approach. In *International Conference on Learning Representations*, 2020.
- [25] K. Zhang, Z. Yang, and T. Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. In *Handbook of Reinforcement Learning and Control*, pages 321–384. 2021.