

Distributed Control using Reinforcement Learning with Temporal-Logic-Based Reward Shaping

Ningyuan Zhang

Wenliang Liu

Calin Belta

Boston University, Massachusetts, USA

NINGYUAN@BU.EDU

WLIU97@BU.EDU

CBELTA@BU.EDU

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We present a computational framework for synthesis of distributed control strategies for a heterogeneous team of robots in a partially observable environment. The goal is to cooperatively satisfy specifications given as Truncated Linear Temporal Logic (TLTL) formulas. Our approach formulates the synthesis problem as a stochastic game and employs a policy graph method to find a control strategy with memory for each agent. We construct the stochastic game on the product between the team transition system and a finite state automaton (FSA) that tracks the satisfaction of the TLTL formula. We use the quantitative semantics of TLTL as the reward of the game, and further reshape it using the FSA to guide and accelerate the learning process. Simulation results demonstrate the efficacy of the proposed solution under demanding task specifications and the effectiveness of reward shaping in significantly accelerating the speed of learning.

Keywords: Multi-Agent Reinforcement Learning, Temporal Logic, Reward Shaping, Partially Observable Stochastic Game

1. Introduction

Many recent works treat the sequential decision-making problem of multiple autonomous, interacting agents as a Multi-Agent Reinforcement Learning (MARL) problem. In this framework, each agent learns a policy to optimize a long-term discounted reward through interacting with the environment and the other agents [Zhang et al. \(2021\)](#); [Qie et al. \(2019\)](#); [Cui et al. \(2019\)](#). In many realistic applications, the environment states are only partially observable to the agents [Hausknecht and Stone \(2015\)](#). Learning in such environments has become a recurrent problem in MARL [Zhang et al. \(2021\)](#), the difficulty of which lies in the fact that the optimal policy might require a complete history of the whole system to determine the next action to perform [Meuleau et al. \(1999\)](#). The authors in [Peshkin et al. \(2001\)](#) used finite policy graphs to represent policies with memory and successfully applied a policy gradient MARL algorithm in a fully cooperative setting.

In this paper, we consider the problem of controlling a heterogeneous team of agents to satisfy a specification given as a Truncated Linear Temporal Logic (TLTL) formula [Li et al. \(2017\)](#) using MARL. TLTL is a version of Linear Temporal Logic (LTL) [Pnueli \(1977\)](#). The semantics of TLTL is given over a finite trajectory in a set such as the state space of the system. TLTL has dual semantics. In its Boolean (qualitative) semantics (whether the formula is satisfied), a trajectory satisfying a TLTL formula is accepted by a Finite State Automaton (FSA). The quantitative semantics assigns a degree of satisfaction (robustness) of a formula. Designing a reward function that accurately

represents the specification is an important issue in reinforcement learning. We use the robustness of the TLTL specification as the reward of the MARL problem. However, the TLTL robustness is evaluated over the entire trajectory, so we can only get the reward at the end of each episode. To address this, we propose a novel reward shaping technique that introduces two additional reward terms based on the quantitative semantics of TLTL and the corresponding FSA. Such rewards, which are obtained after each step, guide and accelerate the learning.

As in [Peshkin et al. \(2001\)](#), we consider a partially observable environment and train a finite policy graph for each agent. When deploying the policy, each agent only knows its own state. However, during training we assume all agents know the states of each other. Then we use this information to guide and accelerate learning. This assumption is reasonable in practice because the training environment is highly configurable and a communication system can be easily constructed, while during deployment there is no guarantee that such a communication system is available.

The idea of modifying rewards using the semantics of temporal logics was introduced in [Li et al. \(2017\)](#) to address the problem of single agent learning in a fully observable environment. The method was then extended to deep RL to successfully control two manipulator arms working together to make and serve hotdogs [Li et al. \(2019\)](#). Although the problem in [Li et al. \(2019\)](#) involved two agents, single agent RL algorithms were used with two different specially designed task specifications. In [Sun et al. \(2020\)](#), temporal logic rewards were used for multi-agent systems. As opposed to our work, the environment in [Sun et al. \(2020\)](#) is fully observable. Moreover, each sub-task in [Sun et al. \(2020\)](#) is restricted to be executed by one agent. In this paper, we allow both independent tasks, which can be accomplished by one agent, and shared tasks, which must be achieved by cooperation of several agents. The rewards in [Li et al. \(2019\)](#) and [Sun et al. \(2020\)](#) encourage the FSA to leave its current state, without distinguishing whether it is a transition towards the satisfaction of the TLTL specification. In this paper, the combination of the two reshaped rewards encourage transitions towards satisfaction. Another related work is [Hammond et al. \(2021\)](#), in which the authors apply an actor-critic algorithm to find a control policy for a multi-agent system that maximizes the probability of satisfying an LTL specification. However, the reward does not capture a quantitative semantics, and a fully observable environment is assumed.

The contributions of this paper can be summarized as follows. First, we propose a general procedure that uses MARL algorithms to synthesize distributed control from arbitrary TLTL specifications for a heterogeneous team of agents. Each agent can have different capabilities and both independent and shared sub-tasks can be defined. The agents work in a partially observable environment, and the policy for each agent only requires its own state during deploying. During training, we assume a fully observable environment, which enables the use of TLTL robustness as a reward. Second, we create a novel temporal logic reward shaping technique, where two additional rewards based on TLTL robustness and FSA states are added. Both rewards are obtained immediately after each step, which guides and accelerates the learning process.

2. Preliminaries and Notation

Given a set S , we use $|S|$, 2^S , and S^* to denote its cardinality, the set of all subsets of S , and a finite sequence over S . Given a set Σ , a collection of sets $\Delta = \{\Sigma_i \subset \Sigma, i \in I\}$, in which I is a set of labels, is called a distribution of Σ if $\cup_{i \in I} \Sigma_i = \Sigma$. For $\sigma \in \Sigma$, we use $I_\sigma = \{i \in I | \sigma \in \Sigma_i\}$ to denote the set of labels of elements in Δ that contains σ .

2.1. Partially Observable Stochastic Games

Single agent reinforcement learning uses Markov Decision Processes (MDP) as mathematical models [Puterman \(2014\)](#). In MARL, more complicated models are needed to describe the interactions between the agents and the environment. A popular model is a Stochastic Game (SG) [Buşoniu et al. \(2010\)](#). A Partially Observable Stochastic Game (POSG), which is a generalization of an SG, is defined as a tuple $\langle I, S, d_0, \{A_i, O_i, B_i\}_{i \in I}, p, \{r_i\}_{i \in I} \rangle$, in which $I = \{1, \dots, n\}$ is an index set for the agents; S is the discrete joint state space; d_0 is the probability distribution over initial states; A_i is the set of actions for agent i , O_i is the discrete observation space and $B_i : S \rightarrow O_i$ is the observation function; $p : S \times A \rightarrow Pr(S)$ is the transition function that maps the states of the game and the joint action of the agents, defined as $A = \times_{i=1}^n A_i$, to probability distributions over states of the game at next time step; and $r_i : S \times A \times S \rightarrow \mathbb{R}$ is the reward function for agent i . If every agent is able to obtain the complete information of the environment state when making decisions, i.e., $O_i = S$ and $B_i(s) = s$ for all $s \in S$, then the POSG becomes a fully observable SG.

The goal for each agent i is to find a policy π_i that maximizes its value V_i :

$$V_i(\boldsymbol{\pi}, d_0) = \sum_{t=0}^{T-1} \gamma^t \mathbb{E}_{a_t \sim \boldsymbol{\pi}} (r_i(s_t, a_t, s_{t+1}) | \boldsymbol{\pi}, d_0),$$

where $\gamma \in (0, 1]$ is the discount factor and $\boldsymbol{\pi} = (\pi_1, \dots, \pi_n)$ is a set of policies. Note that for each agent the environment is non-stationary since its reward r_i may depend on other agents' policies. A *memory-less policy* $\pi_i : O_i \rightarrow Pr(A_i)$ is a mapping from the observations of agent i to probability distributions over the action space of agent i . For a fully observable SG, memory-less policies are sufficient to achieve optimal performance. However, for POSGs, the best memoryless policy can still be arbitrarily worse than the best policy using memory [Singh et al. \(1994\)](#). A policy graph [Meuleau et al. \(1999\)](#) is a common way to represent a policy with memory.

2.2. Truncated Linear Temporal Logic

Truncated Linear Temporal Logic (TLTL) [Li et al. \(2017\)](#) is a predicate temporal logic inspired from the traditional Linear Temporal Logic (LTL) [Pnueli \(1977\)](#). TLTL can express rich task specifications that are satisfied in finite time. Its formulas are defined over predicates $f(x) \geq 0$, where $f : X \rightarrow \mathbb{R}$ is a function over X (such as the state space of a system). Then TLTL formulas are evaluated against finite sequences over X . TLTL formulas have the following syntax:

$$\phi := \top \mid f(x) \geq 0 \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \Rightarrow \psi \mid \diamond\phi \mid \bigcirc\phi \mid \phi \mathcal{U} \psi \mid \phi \mathcal{T} \psi,$$

where \top is the Boolean True. \neg (negation), \wedge (conjunction) and \vee (disjunction) are Boolean connectives. \diamond , \bigcirc , \mathcal{U} , and \mathcal{T} are temporal operators that stand for “eventually”, “next”, “until” and “then” respectively. Given a sequence over X , $\diamond\phi$ (eventually) requires ϕ to be satisfied at some time step, $\bigcirc\phi$ (next) requires ϕ to be satisfied at the second time step, $\phi \mathcal{U} \psi$ (until) requires ϕ to be satisfied at each time step before ψ is satisfied, $\phi \mathcal{T} \psi$ (then) requires ϕ to be satisfied at least once before ψ is satisfied. TLTL also has derived operators, e.g., \square (finite time always) and \Rightarrow (imply).

TLTL formulas can be interpreted in a qualitative semantics or a quantitative semantics. The qualitative semantics provides a yes/no answer to the corresponding property, while the quantitative semantics generates a measure of the degree of satisfaction. Given a finite trajectory $x_{t:t+k} := x_t x_{t+1} \dots x_{t+k}$, the quantitative semantics of a formula ϕ , also called *robustness*, is denoted by

$\rho(x_{t:t+k}, \phi)$. We refer the readers to [Li et al. \(2017\)](#) for detailed definitions of the TLTL qualitative and quantitative semantics.

2.3. Finite State Automata

Definition 1 (Finite State Predicate Automaton [Li et al. \(2019\)](#)) *A finite state predicate automaton (FSPA) corresponding to a TLTL formula ϕ is a tuple $\mathcal{A} = \langle Q, q_0, F, T_r, \Psi, \mathcal{E}, b \rangle$, in which Q is a set of automaton states, $q_0 \in Q$ is the initial state, $F \subset Q$ is the set of final states, T_r is the set of trap states, Ψ is a set of predicate Boolean formulas, where the predicates are evaluated over X , $\mathcal{E} \subseteq Q \times Q$ is the set of transitions, and $b : \mathcal{E} \rightarrow \Psi$ maps transitions to formulas in Ψ .*

The semantics of FSPA is defined over finite sequences over X . We refer to the elements in Ψ as *edge guards*. A transition is enabled if the corresponding Boolean formula is satisfied. Formally, by noting that a Boolean formula is a particular case of a TLTL formula, the FSPA transitions from q_t to q_{t+1} at time t if and only if $\rho(x_{t:t+k}, b(q_t, q_{t+1})) > 0$, where $x_{t:t+k}$ is a sequence of X from t to $t+k$. Note that the robustness is only evaluated at s_t instead of the whole sequence. Thus we abbreviate $\rho(x_{t:t+k}, b(q_t, q_{t+1}))$ to $\rho(x_t, b(q_t, q_{t+1}))$.

A trajectory over Q is accepting if it ends in the final states F . A trajectory over X is accepted by \mathcal{A} if it leads to an accepting trajectory over Q . Each TLTL can be translated into an equivalent FSPA [Li et al. \(2019\)](#), in the sense that a trajectory over X satisfies the TLTL specification if and only if the same trajectory is accepted by the FSPA. All trajectories over X that violate the TLTL formula drive the FSPA to the trap states in T_r .

3. Problem Formulation

The geometry of the world is modeled by an environment graph $G = (V, E)$, where V is a set of vertices and $E \subset V \times V$ is a set of edges. The motions of the agents are restricted by this graph. Assume we have a set of agents $\{i | i \in I\}$ where I is a label set, and a set of service requests Σ . Let $l : \Sigma \rightarrow 2^V$ be a function indicating the locations at which a request occurs. For a given request $\sigma \in \Sigma$, let $V_\sigma = l(\sigma) \subseteq V$ be the set of vertices where σ occurs. The definition of V_σ implicitly assumes that one request can occur at multiple vertices, since $V_\sigma \in 2^V$. Also, multiple requests can occur at the same vertex, since V_{σ_1} and V_{σ_2} may have non-empty intersection for $\sigma_1 \neq \sigma_2$. We use a distribution $\Delta = \{\Sigma_i \subset \Sigma, i \in I\}$ to model the agents' capabilities for completing different service requests: $\sigma \in \Sigma_i$ means that request σ can be serviced by agent i . We further define $I_\sigma = \{i \in I | \sigma \in \Sigma_i\}$ as the set of agents that can service request σ . We consider two types of services: *independent* and *shared*. An independent request is any $\sigma \in \Sigma$ such that $|I_\sigma| = 1$. This type of request can only be serviced by agent i such that $\sigma \in \Sigma_i$. A shared request σ is defined by $|I_\sigma| > 1$, which means that servicing it requires the cooperation of all the agents that own σ .

We model the motion and actions of agent i by using a deterministic finite transition system $\mathcal{T}_i = \langle V_i, v_{0i}, A_i, \delta_i, \Sigma_i, \models_i \rangle$ where $V_i \subseteq V$ is the set of vertices that can be reached by agent i ; $v_{0i} \in V_i$ is the initial location; $A_i = V_i \cup \Sigma_i \cup \{\epsilon\}$ is the discrete action space; $\delta_i : V_i \times A_i \rightarrow V_i$ is a deterministic transition function; $\Pi_i = \Sigma_i \cup \{\epsilon\}$ is a proposition set; $\models_i \subseteq V \times \Pi_i$ is the satisfaction relation such that 1) $(v, \epsilon) \in \models_i$ for all $i \in I$, $v \in V_i$, and 2) $(v, \sigma) \in \models_i$, $\sigma \in \Sigma_i$, if and only if $v \in l(\sigma)$. The meaning of taking action $a \in A_i$ at state v is as follows: if $a \in V_i$, then the agent tries to move to vertex a from v ; if $a \in \Sigma_i$, then agent i tries to conduct service a at v ; $a = \epsilon$ indicates

that the agent stays at the current vertex without conducting any requests. Formally,

$$\delta_i(v, a) = \begin{cases} v' & \text{if } a = v' \in V_i \text{ and } (v, v') \in E \\ v & \text{otherwise} \end{cases} \quad (1)$$

For $v, v' \in V_i$ and $a \in A_i$, a transition $v' = \delta_i(v, a)$ is also denoted by $v \xrightarrow{a}_{\mathcal{T}_i} v'$. Now we define the Motion and Service (MS) plan for an agent, which is inspired by [Chen et al. \(2011\)](#).

Definition 2 (Motion and Service Plan) *A Motion and Service (MS) plan for an agent $i \in I$ is a sequence of states represented by ordered pairs $z_0^i z_1^i \dots z_T^i \in (V \times \Pi_i)^*$ that satisfies the following properties:*

1. $z_0^i = (v_{0i}, \epsilon)$.
2. For all $t \geq 0$, $z_t^i = (v, \sigma) \in \models_i$.
3. For all $t \geq 1$, given $z_{t-1}^i = (v, \sigma)$ and $z_t^i = (v', \sigma')$, if $v \neq v'$ then $\sigma' = \epsilon$.
4. For all $t \geq 1$, given $z_{t-1}^i = (v, \sigma)$ and $z_t^i = (v', \sigma')$, then $\exists a \in A_i$ such that $\delta_i(v, a) = v'$.

We use $z_{0:T}^i$ to denote the MS plan for an agent i from time 0 to T , i.e., $z_{0:T}^i = z_0^i z_1^i \dots z_T^i$. A Team Motion and Service (TMS) plan is then defined as $z_{0:T} = \times_{i=1}^n z_{0:T}^i$.

We assume that there exists a global discrete clock that synchronizes the motions and the services of requests of all agents. We also assume that the times needed to service requests are all equal to 1. In other words, similar to POSGs, at each time step, an agent either chooses to move to a vertex according to δ_i or to stay where it is to conduct a particular request in Σ_i . Before the beginning of the next time step, all motions and requests are completed and the agents are ready to execute the next state in their MS plans. A MS plan thus uniquely defines a sequence of actions for an agent. More specifically, the action derived from a MS plan z^i at $t \geq 1$ is determined by $z_{t-1}^i = (v, \sigma)$ and $z_t^i = (v', \sigma')$. If $\sigma' = \epsilon$ and $v \neq v'$, then the agent moves to vertex v from v' (for the case where $v = v'$, the agent does not conduct any request and waits at v). On the other hand, $\sigma \neq \epsilon$ means that the agent should conduct request σ and it must have reached vertex v at the previous time step according to property 3 from the above definition.

In this paper, we use the term *global behavior* to refer to the sequence of requests serviced by the whole team. An independent request σ is considered to be finished at time t if and only if $z_t^i = (v, \sigma)$, $i \in I_\sigma$. For a shared request σ , we assume that all the agents that own this request are capable of communicating with each other at the same vertex where σ occurs. A shared service request σ is completed at time t if and only if $z_t^i = (v_i, \sigma)$ for all $i \in I_\sigma$ and $v_j = v_i$ for all $i, j \in I_\sigma, i \neq j$. Thus, given individual MS plans of all the agents, one single global behavior is then uniquely determined. This is directly deduced from the assumption of a constant finishing time for all requests. Next we define a term called *Team Trajectory* to describe the results of executing a global TMS plan.

Definition 3 (Team Trajectory) *Given a team of n agents \mathcal{T}_i , a set of service requests Σ , a distribution $\Delta = \{\Sigma_i \subseteq \Sigma, i \in I\}$, a function $l : \Sigma \rightarrow V$ that shows the locations of requests and a TMS plan $z_{0:T}$ for the whole team of agents, the Team Trajectory from executing the TMS plan is a sequence of $T + 1$ states $x_{0:T}$, each of which is a $(n + 1)$ -tuple, with the first n elements equal to the vertices in the corresponding individual MS plan. The last element of the team trajectory at each time t is a set $\bar{\Sigma}_t \subseteq \Sigma$, such that $\bar{\Sigma}_t := \{\sigma \in \Sigma \mid \sigma \text{ is finished at time } t\}$.*

Now the problem considered in this paper can be formulated as follows.

Problem 3.1 *Given an environment graph G , a team of agents \mathcal{T}_i as defined above, $i \in I$, a set of service requests Σ , a distribution $\Delta = \{\Sigma_i \subseteq \Sigma, i \in I\}$, a function $l : \Sigma \rightarrow V$ that shows locations of requests, find a set of MS plans for each agent such that the team trajectory obtained by executing individual MS plans satisfying a given global task specification encoded by a TLTL formula over predicate functions of states of the team trajectory $x_{0:T}$.*

A natural, classical motion planning approach to Problem 3.1 would be to design in advance MS plans for each agent. However, this requires the transition function δ_i for each agent i to be known, which is not always true in realistic applications. Moreover, a priori top-down design that guarantees satisfactory performance can become extremely difficult in complex and time-varying environments [Buşoniu et al. \(2010\)](#). Therefore, we tackle this problem using MARL, in which agents learn how to act by constantly interacting with the environment and other agents, and adjusting their behaviors according to feedback received from the environment. The transition functions δ_i are assumed to be unknown.

Example 1 *Consider a team of two heterogeneous agents that have to service three requests: σ_1, σ_2 and σ_3 , with an additional requirement that σ_3 must not be finished until σ_1 or σ_2 has been finished. The environment is shown in Fig. 3 a 5×5 as a grid world, which can be abstracted by a graph $G = (V, E)$. Each cell represents a vertex $v \in V$ (labeled with values at the bottom right) and each facet forms a reflective, two-way edge in E . There are three service requests: $\Sigma = \{\sigma_1, \sigma_2, \sigma_3\}$. There is a team of two robots modeled by transition systems $\mathcal{T}_1 = \langle V, v_4, A_1, \delta_1, \Sigma_1, \models_1 \rangle$ and $\mathcal{T}_2 = \langle V, v_{18}, A_2, \delta_2, \Sigma_2, \models_2 \rangle$. Assume both agents are able to move through any edges in E . Thus the action spaces are $A_1 = V \cup \Sigma_1 \cup \{\epsilon\}$ and $A_2 = V \cup \Sigma_2 \cup \{\epsilon\}$. The capabilities of conducting services are captured by the distribution $\Delta = \{\Sigma_1, \Sigma_2\}$, $\Sigma_1 = \{\sigma_1, \sigma_3\}$ and $\Sigma_2 = \{\sigma_2, \sigma_3\}$, which means σ_1 and σ_2 are independent requests and σ_3 is a shared request. For instance, to finish request σ_1 , robot \mathcal{T}_1 must move to vertex $v_{21} \in l(\sigma_1)$ and then choose to take action $\sigma_1 \in A_1$. This behavior, which can be interpreted as a sub-task towards the success of the global mission, is formulated by a TLTL formula $\phi_1 = go_1(\cdot) \mathcal{T} do_1(\cdot) \wedge \Box \neg (\neg go_1(\cdot) \wedge do_1(\cdot))$ where $go_1(\cdot)$ and $do_1(\cdot)$ are predicate functions defined over states of the team trajectory as*

$$\begin{aligned} go_1(x_{t:T}) &= C_1 - \min_{v \in l(\sigma_1)} dist(v_t^1, v) \\ do_1(x_{t:T}) &= \begin{cases} C_2 & \sigma_1 \in \bar{\Sigma}_t \\ -C_2 & \text{otherwise} \end{cases} \end{aligned} \quad (2)$$

where v_t^1 is the location of robot 1 at time t (included in team trajectory state x_t), $dist(v_t^1, l(\sigma_1))$ is the distance from v_t^1 to the location at which σ_1 occurs, and C_1 and C_2 are positive constants. Then the global specification is given by the TLTL formula $\phi = (\Diamond \phi_1 \vee \Diamond \phi_2) \wedge \Diamond \phi_3 \wedge \neg \phi_3 \mathcal{U} (\phi_1 \vee \phi_2)$. The FSPA corresponding to ϕ can be found in Fig. 3. The predicate functions $go_2(\cdot)$, $do_2(\cdot)$ and $do_3(\cdot)$ are defined similarly as (2). Predicates $go_3(\cdot)$ are designed in a slightly different form to encode the cooperative behaviors:

$$go_3(x_{t:T}) = C_3 - \min_{v \in l(\sigma_3)} \left(\max_{i \in I_{\sigma_3}} [dist(v_t^i, v)] \right), \quad (3)$$

where C_3 is some constant. Note that although formula ϕ is constructed hierarchically, the resultant FSPA is actually nonhierarchical.

where $\psi_{q_t, q_{t+1}} = b(q_t, q_{t+1}) \in \Psi$ is the edge guard of the transition from q_t to q_{t+1} and x_t is a state of the Team Trajectory included in s_t . Let the reward function $r_1 = \dots = r_n = r$ be defined as

$$r(s_t, a_t, s_{t+1}) = \begin{cases} C & \text{if } q_{t+1} \in F \text{ and } q_t \notin F \\ -C & \text{if } q_{t+1} \in T_r \text{ and } q_t \notin T_r \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $C > 0$ is a constant.

In this paper, we assume that the FSPA-POSG has a finite horizon T . The rationale of this assumption lies in the fact that any accepting or rejecting team trajectory must be finite (the corresponding FSPA reaches one state either in F or in T_r). However, T must be large enough to give agents adequate time to complete the task and thus it becomes a design factor of our approach. Clearly, since the team trajectory state is included in the FSPA-POSG state, constructing a team trajectory $x_{0:T}$ from $s_{0:T}$ is very simple, then constructing a TMS is also straightforward.

4.2. Temporal Logic Reward Shaping

As discussed in Sec. 1, we assume that the state of the FSPA-POSG is fully observable during training, while the observation function during deploying is as defined in Def. 4. The reward function defined above only provides feedback to agents at the very end of each episode, which could cause serious problems when the probability of success under randomly selected actions converges to zero as the mission becomes more and more complicated. Missions described by temporal logics are expected to be complicated. Therefore, reward shaping techniques are introduced to add intermediate rewards to improve the speed and successful rate of learning.

We developed two additional rewards for each agent i to guide learning: r_ρ^i and r_J^i . r_ρ^i is defined based on TLTL's quantitative semantics as follows:

$$r_\rho^i(s_t, a_t, s_{t+1}) = \rho(x_{t+1}, D_{q_{t+1}}) - \rho(\hat{x}_{t+1}^i, D_{\hat{q}_{t+1}}), \quad (5)$$

where $D_{q_{t+1}} = \bigvee_{(q_{t+1}, q_{t+2}) \in \mathcal{E}, q_{t+2} \notin T_r} b(q_{t+1}, q_{t+2})$ is the disjunction of all predicates of the outgoing edges from q_{t+1} to other non-trapped states, $\hat{x}_{t+1}^i = (v_t^i \times_{j=1, j \neq i}^n v_{t+1}^j, \hat{\Sigma}_{t+1})$ is a predicted FSPA-POSG state that all agents take actions and transition to the next state except agent i , \hat{q}_{t+1} is the corresponding predicted automaton state. The robustness value $\rho(\cdot, D_{q_{t+1}})$ can be interpreted as a metric that measures how far it is from leaving the FSPA state q_{t+1} given x_{t+1} . The individual reward r_ρ^i is designed to be the difference between the robustness value under the whole team's actions and the robustness value with agent i being idle. This is reasonable because of two reasons. First, \hat{x}_{t+1}^i is generated by only using past observations so it does not require the agent model δ_i . Second, since we know Σ_i , the value of predicate functions $do_i(\cdot)$ can be predicted by removing all the services that requires agent i 's action from the set of services that are finished at t , which makes \hat{q}_{t+1} available. r_ρ^i is able to distinguish each agent's contribution to the global behavior, and thus provide more informative rewards to agents.

By incorporating r_ρ^i , the agents are encouraged to leave the current FSPA state. However, the limitation of r_ρ^i is that it cannot distinguish whether a transition is beneficial or not. For example, in Fig. 3, the bad transitions that we want to avoid are colored by orange. As for state q_5 , two agents have already reached the location of σ_3 and the next thing to do is to finish σ_3 cooperatively. However, if they leave the vertex where σ_3 lives, r_ρ^i will generate positive rewards because the FSPA

has left q_5 , which clearly violates our intentions. It is true that the learning can still converge to the correct policies, but only through the effect of discount factor $\gamma < 1$, in the sense that transition from q_5 to q_4 will delay the time of completing the mission and thus yield a lower final reward.

To overcome this issue, we design another reshaped reward signal r_J^i , which can be seen as a "potential energy" of the corresponding FSPA state. Before giving the definition of r_J^i , we first introduce the concept of path length. Let $H(q, q')$ denote the set of all finite trajectories from $q \in Q$ to $q' \in Q$. The path length L is defined as

$$L(\mathbf{q}) = \sum_{k=1}^{n-1} \omega_{\mathcal{A}}(q_k, q_{k+1}), \quad (6)$$

where $\mathbf{q} = q_1 \dots q_n \in H(q_1, q_n)$ is a finite sequence of FSPA states connected by transitions in \mathcal{E} and $\omega_{\mathcal{A}} : \mathcal{E} \rightarrow \mathbb{R}^+$ is a weight function. Then a distance function from q to q' can be defined as

$$d(q, q') = \begin{cases} \min_{\mathbf{q} \in H(q, q')} L(\mathbf{q}) & \text{if } H(q, q') \neq \emptyset \\ \infty & \text{if } H(q, q') = \emptyset. \end{cases} \quad (7)$$

Now an energy function $J(q)$, $q \in Q$ can be defined as

$$J(q) = \begin{cases} 0 & \text{if } q \in F \\ \min_{q' \in F} d(q, q') & \text{if } q \notin F. \end{cases} \quad (8)$$

In words, the energy function J of a state q is the minimum weighted sum of transitions it takes to reach the set of final states. Finally, r_J^i is constructed as follows:

$$r_J^1 = \dots = r_J^n = J(q) - J(q_{t+1}). \quad (9)$$

r_J^i ensures that agents will receive an assessment of every FSPA transition immediately. There is no standard procedure for determining the weight function $\omega_{\mathcal{A}}$ and hence it becomes a design factor.

4.3. Solving the equivalent FSPA-POSG

We use the policy gradient method with policy graphs as in [Meuleau et al. \(1999\)](#) and [Peshkin et al. \(2001\)](#) to find the optimal policy for each agent in the FSPA-POSG. In the training phase, both the original reward and the two reshaped rewards require the information of the FSPA state, which is determined by all the agents. Hence, a centralized coordinator is necessary. We assume that the communication among all agents during the training phase is available, which is reasonable as stated in [Sec. 1](#). After training, each agent obtains a policy that tracks its own history states to decide the action to be taken at each time step, i.e., the policies are distributed.

5. Simulation

To verify the efficacy of the method described in [Sec. 4](#), we conducted 500 simulations of training on [Example 1](#) with reshaped rewards r_{ρ}^i , $r_{\rho}^i + r_J^i$ and 0 (only original reward), respectively. In each simulation, we trained the policies for 10000 episodes. We set the discount factor γ to 0.995, the learning step size α to 0.01, the constant C in [Eqn. \(4\)](#) to 1. The number of internal states of the policy graphs are chosen to be 10 for both agents. In the process of training policy graphs, we utilized Boltzmann exploration with a temperature parameter of 1. As a comparison, we also




Figure 3: The averaged learning curves of the three rewards. The values on the y -axis represent the normalized original reward signal in Eqn. (4).




Figure 4: Final trajectories generated by policy graphs after training.

Table 1: Rates of convergence to the optimal policies with different rewards

Reshaped rewards	r_ρ^i	$r_\rho^i + r_J^i$	0	$r_\rho^i + r_J^i$ (memoryless)
Optimal policy convergence rate	68.4%	83.0%	70.2%	0%

applied decentralized Q-learning [Matignon et al. \(2012\)](#) with reshaped rewards $r_\rho^i + r_J^i$ and an ϵ -greedy exploration strategy to learn policies without memory.

It was not guaranteed that agents’ policies would always converge to the optimal ones. The rates at which the optimal policies were learned within 10000 episodes are shown in Table 1. We can see that using $r_\rho^i + r_J^i$ boosted the convergence rate by 12.8% compared to the original reward. Note that memoryless Q-learning completely fails the task. The averaged learning curves in which optimal policies successfully converged are shown in Fig. 3. We can see that the proposed reward shaping technique can accelerate the speed of learning and significantly reduce the variance.

The behaviors of agents governed by the learned policies are visualized in Fig. 4. In example 1, there are two possible routes to finish the TLTL task: finish σ_1 then σ_3 ; finish σ_2 then σ_3 . It can be seen that the agents are able to figure out the optimal policy with minimum route length.

We also inspected the final policy graphs and their corresponding trajectories in simulations where agents with reshaped rewards failed to satisfy the TLTL specification. It turns out that in the majority of the failures, the agents would first finish either σ_1 or σ_2 and then never reach σ_3 by wondering among several vertices near σ_1 or σ_2 . One possible cause is that the gradient ascent took steps that were too large, which cannot be fully handled by simply decreasing α [Schulman et al. \(2017\)](#). More advanced policy gradient methods such as PPO [Schulman et al. \(2017\)](#) or TRPO [Schulman et al. \(2015\)](#) might solve this issue. We will explore this direction in future research.

6. Conclusion

In this paper, we applied MARL to synthesize distributed controls for a heterogeneous team of agents from Truncated Linear Temporal Logic (TLTL) specifications. We assumed a partially observable environment, where each agent’s policy depends on the history of its own states. We used the TLTL robustness as the reward of the MARL and introduce two additional reshaped rewards. Simulation results demonstrated the efficacy of our framework and showed that reward shaping significantly improves the convergence rate to the optimal policy and the learning speed. Future work includes consideration of non-deterministic SGs and other policy gradient methods.

Acknowledgments

This work was supported by NSF under Grant IIS-2024606 and Grant OIA-2020983.

References

- Lucian Buşoniu, Robert Babuška, and Bart De Schutter. Multi-agent reinforcement learning: An overview. *Innovations in multi-agent systems and applications-1*, pages 183–221, 2010.
- Yushan Chen, Xu Chu Ding, Alin Stefanescu, and Calin Belta. Formal approach to the deployment of distributed robotic teams. *IEEE Transactions on Robotics*, 28(1):158–171, 2011.
- Jingjing Cui, Yuanwei Liu, and Arumugam Nallanathan. Multi-agent reinforcement learning-based resource allocation for uav networks. *IEEE Transactions on Wireless Communications*, 19(2): 729–743, 2019.
- Lewis Hammond, Alessandro Abate, Julian Gutierrez, and Michael Wooldridge. Multi-agent reinforcement learning with temporal logic specifications. *arXiv preprint arXiv:2102.00582*, 2021.
- Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. In *2015 aaai fall symposium series*, 2015.
- Xiao Li, Cristian-Ioan Vasile, and Calin Belta. Reinforcement learning with temporal logic rewards. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3834–3839. IEEE, 2017.
- Xiao Li, Zachary Serlin, Guang Yang, and Calin Belta. A formal methods approach to interpretable reinforcement learning for robotic planning. *Science Robotics*, 4(37), 2019.
- Laetitia Matignon, Guillaume J Laurent, and Nadine Le Fort-Piat. Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. *The Knowledge Engineering Review*, 27(1):1–31, 2012.
- Nicolas Meuleau, Leonid Peshkin, Kee-Eung Kim, and Leslie Pack Kaelbling. Learning finite-state controllers for partially observable environments. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 427–436, 1999.
- Leonid Peshkin, Kee-Eung Kim, Nicolas Meuleau, and Leslie Pack Kaelbling. Learning to cooperate via policy search. *arXiv preprint cs/0105032*, 2001.
- Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 46–57. IEEE, 1977.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Han Qie, Dianxi Shi, Tianlong Shen, Xinhai Xu, Yuan Li, and Liujing Wang. Joint optimization of multi-uav target assignment and path planning based on multi-agent reinforcement learning. *IEEE access*, 7:146264–146272, 2019.

- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Satinder P Singh, Tommi Jaakkola, and Michael I Jordan. Learning without state-estimation in partially observable markovian decision processes. In *Machine Learning Proceedings 1994*, pages 284–292. Elsevier, 1994.
- Chuangchuang Sun, Xiao Li, and Calin Belta. Automata guided semi-decentralized multi-agent reinforcement learning. In *2020 American Control Conference (ACC)*, pages 3900–3905. IEEE, 2020.
- Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of Reinforcement Learning and Control*, pages 321–384, 2021.