

Learning Symbolic Task Decompositions for Multi-Agent Teams

Ameesh Shah*
UC Berkeley
Berkeley, USA
ameesh@berkeley.edu

Niklas Lauffer*
UC Berkeley
Berkeley, USA
nlauffer@berkeley.edu

Thomas Chen*
UC Berkeley
Berkeley, USA
thomasychen@berkeley.edu

Nikhil Pitta*
UC Berkeley
Berkeley, USA
nikhil.pitta@berkeley.edu

Sanjit A. Seshia
UC Berkeley
Berkeley, USA
sseshia@berkeley.edu

ABSTRACT

One approach for improving sample efficiency in cooperative multi-agent learning is to decompose overall tasks into sub-tasks that can be assigned to individual agents. We study this problem in the context of *reward machines*: symbolic tasks that can be formally decomposed into sub-tasks. In order to handle settings without *a priori* knowledge of the environment, we introduce a framework that can learn the optimal decomposition from model-free interactions with the environment. Our method uses a task-conditioned architecture to simultaneously learn an optimal decomposition and the corresponding agents’ policies for each sub-task. In doing so, we remove the need for a human to manually design the optimal decomposition while maintaining the sample-efficiency benefits of improved credit assignment. We provide experimental results in several deep reinforcement learning settings, demonstrating the efficacy of our approach. Our results indicate that our approach succeeds even in environments with codependent agent dynamics, enabling synchronous multi-agent learning not achievable in previous works.¹

KEYWORDS

Decentralized Multi-Agent Learning, Discrete Event Systems, Multi-Agent Reinforcement Learning, Reward Machines

ACM Reference Format:

Ameesh Shah, Niklas Lauffer, Thomas Chen, Nikhil Pitta, and Sanjit A. Seshia. 2025. Learning Symbolic Task Decompositions for Multi-Agent Teams. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 10 pages.

1 INTRODUCTION

Using a single reward signal for a team of agents in multi-agent reinforcement learning (MARL) can make it challenging for agents to understand how their individual behavior impacts the overall reward. This challenge in MARL is known as the *credit assignment*

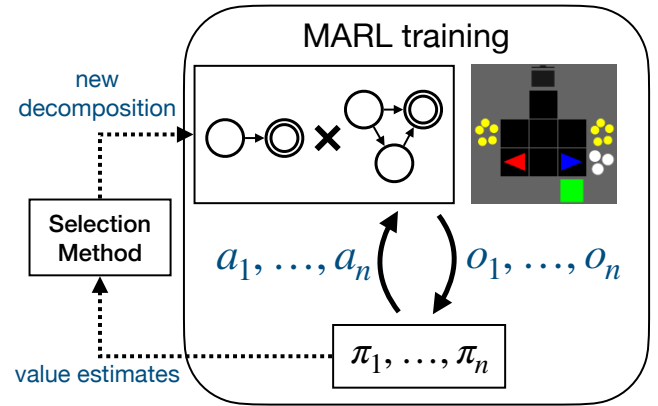


Figure 1: Visualization of our learning framework. At each new episode of training, a selection method chooses a possible symbolic decomposition of the task and assigns sub-tasks to each agent in the team. As each agent learns the viability of different sub-tasks, our selection method simultaneously finds the optimal task decomposition.

problem [1] and can severely limit the effectiveness of naive reinforcement learning approaches in the multi-agent setting [23, 36].

One method for addressing the credit assignment problem is to formulate the task as a *symbolic concept* that can be precisely decomposed into sub-tasks for assignment to individual agents [21, 34]. By using the reward signal from each sub-task, agents are credited for completing the specific sub-task they are assigned, even if the overall task is not achieved.

Previous literature has established sufficient conditions for “valid” decompositions of multi-agent *reward machines* [38], an automaton-based symbolic task structure, where satisfaction of the sub-tasks provably satisfies the overall task [21]. Human-designed decompositions that satisfy these validity conditions help the agents learn to accomplish the task more quickly in tabular settings. Further work demonstrates that valid decompositions can be automatically generated based on human-designed heuristics [34].

These prior approaches to reward machine decomposition are limited when many possible decompositions exist. In these cases, the optimal decomposition of a task must be carefully designed and selected by a human. Selecting a meaningful task decomposition amongst many possible options often requires prior knowledge

*Equal Contribution

¹Code is available at <https://github.com/thomasychen/LOTaD>



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

of the environment that is not assumed in standard RL, making the process tedious at best and impossible at worst. If a task decomposition is selected arbitrarily, the decomposition may not be compatible with the specifics of the environment to effectively break down the task.

Motivating example. Consider the “Repairs” environment in Figure 2, where a team of three robot agents must visit a number of communication stations in their environment to make repairs. First, any two of the agents must meet at headquarters (HQ, denoted by the control tower symbol), at which point a ready signal is sent out to both stations (red and yellow) to inform the stations that agents will be visiting each station to perform repairs. The agents are then tasked to visit these stations in any order. The agents can independently move around in grid in any of the four cardinal directions, or remain still. We can formulate this task in the form of a *Reward Machine* (RM), an automaton that encodes the objective over high level ‘events’ which occur in specific states of the environment (visualized in Figure 2). RMs offer a precise notion of task completion and can easily capture temporal dependencies required in objectives (i.e., visit the stations *after* HQ).

In the Repairs environment, there exists a hazardous region, encoded in orange, that prevents more than one agent from entering the region at a time. The location and existence of this hazardous region is *not* known to the agents a priori, which prevents this constraint from being encoded in the reward machine as part of the task description.

This task, even with no prior knowledge of the environment, can naturally be decomposed in a number of ways. For example, Agent 1 and Agent 2 can be tasked with meeting at HQ, and then Agent 2 can visit the yellow station and leave the hazardous region while Agent 3 visits the red station. Alternatively, Agents 2 and 3 can meet at HQ, and then Agent 1 can visit both the red and yellow stations. Although many plausible decompositions exist, knowing which decomposition of the overall task leads to the most efficient completion of the task (i.e., is *optimal*) largely depends on the dynamics of the underlying environment. In our case, Agent 3 happens to be closest to the red station, which might mean that assigning that station to them would be optimal. However, without knowledge of the layout of the environment, as is standard in a model-free setting, this would be impossible to know ahead of time.

Our Contributions. In this work, we address the aforementioned limitations by introducing *an approach to automatically find optimal decompositions of an overall task into sub-tasks* for a team of agents. Our method is a lightweight extension of standard MARL and is applicable to model-free settings with no prior information about the environment or individual capabilities of the agents.

We summarize our approach briefly as follows. At the beginning of training, we generate possible *candidate* decompositions of our overall task as assignments of sub-tasks for our agents to accomplish. Then, during training, we explore selecting different decompositions for our team of agents and observe their performance as they attempt to learn a goal-conditioned policy that can achieve the variety of possible sub-tasks captured by our candidates. We record the value of agents’ performances for different sub-tasks, and use this information to intelligently select subsequent decompositions for our agent team. This allows us to simultaneously learn (1) the optimal decomposition for our task and (2) policies for each agent




Figure 2: (Top) The “Repairs” MDP with a team of 3 agents. (Bottom) A task completion reward machine (RM) encoding the task: agents must navigate the environment to visit the HQ control tower, and then visit a set of communication stations. The goal state of the RM is denoted by concentric circles.

that optimize said decomposition. We leverage selection strategies popularized in the multi-armed bandit algorithm literature [4, 6, 17] to balance exploring different candidate decompositions with exploiting the value our agents achieve as they learn sub-tasks during training. We visualize our approach in Figure 1.

In addition to our decomposition selection strategy, we introduce a novel training setup for RM-driven MARL that rectifies issues caused by dependent agent dynamics. Prior approaches [21, 34] to RM-driven MARL make the assumption that the dynamics of each agent are independent of other agents, so that each agent can be trained individually in an environment absent from other agents. This assumption is often impractical: it is inefficient to train individual agents independently, and in many cases, multi-agent dynamics are codependent. For example, an agent that has completed their sub-task may obstruct another agent from the completion of their own sub-task. Our setup gives each agent in the team a *global* view of the overall task along with their sub-task, enabling agents to learn policies that accomplish their individual sub-tasks and help facilitate the completion of other agent’s sub-tasks as well.

We summarize our contributions as follows: (1) We introduce a method for learning optimal decompositions from model-free interactions with the environment. (2) Within our method, we provide a novel training architecture that allows multiple agents to train simultaneously within an environment and avoids conflicts that arise due to dependent dynamics across agents. (3) We demonstrate our approach’s improvement in learning multi-agent policies against baseline approaches in several cooperative environments. Our results show that MARL benefits substantially from the improved credit assignment of task decompositions even when the optimal decomposition is not known *a priori*.

2 PRELIMINARIES

2.1 MDPs and Labeled Markov Games

A Markov Decision Process (MDP) $\mathcal{M} = (S, A, T^{\mathcal{M}}, d_0, \gamma)$ is a tuple that consists of a state space S , an action space A , a transition function $T^{\mathcal{M}} : S \times A \rightarrow S$, an initial state distribution $d_0 \in \Delta(S)$, and a discount factor $0 < \gamma < 1$. A (stationary) policy $\pi : S \rightarrow \Delta(A)$ in \mathcal{M} produces a distribution over actions given a state. Following standard RL notation, we define $\pi(a_t | s_t)$ as the probability of taking action a_t in state s_t at timestep t . A policy takes an action in a given state in \mathcal{M} , transitions to a new state via $T^{\mathcal{M}}$, and repeats, generating a **trajectory** of length τ as $(s_0, a_0, \dots, s_\tau, a_\tau)$.

To generalize to the MARL setting, we extend \mathcal{M} to a cooperative Markov game with homogeneous action spaces. We define a cooperative Markov game with n agents as $\mathcal{G} = (S, \mathcal{A}, T^{\mathcal{G}}, d_{01} \dots d_{0n}, \gamma, L^{\mathcal{G}})$, which corresponds to the joint set of states $S = S_1 \times \dots \times S_n$, the joint set of actions $\mathcal{A} = A_1 \times \dots \times A_n$, a joint transition operator $T^{\mathcal{G}} : S \times \mathcal{A} \rightarrow S$, a set of independent initial distributions for each agent $d_{01} \times \dots \times d_{0n}$, and a discount factor γ that remains unchanged. We consider the *centralized training, decentralized execution* [7] setting of MARL, where each agent receives independent observations and deploys their individual policies $\pi_i : S_i \rightarrow \Delta(A_i)$ at execution time, but are jointly trained. The joint policy $\pi : S \rightarrow \Delta(\mathcal{A})$ executes all individual policies simultaneously at each timestep, and successive states are dictated by $T^{\mathcal{G}}$, generating joint trajectories $((s_0^0, a_0^0, \dots, s_0^n, a_0^n), \dots, (s_\tau^0, a_\tau^0, \dots, s_\tau^n, a_\tau^n))$.

In addition to the standard components of our Markov Game, we assume access to a known *labeling function* $L^{\mathcal{G}} : S \rightarrow 2^\Sigma$ that maps states in S to a set of *environment events*, denoted by Σ . Each environment event $e \in \Sigma$ is represented by a variable that takes on a Boolean truth value (True or False). For example, in Figure 2, if Agent 1 is in the yellow station, Agent 2 is in the red station, and Agent 3 has not moved from its initial position, the labeling function for this joint state would return $\{Y_S, R_S\}$.

2.2 Task completion Reward Machines

In this work, we consider task completion reward machines (RMs) [21, 34, 45] as our team objective. A task completion RM is specified by a tuple $\mathcal{R} = (U, u_{-1}, \Sigma, \delta, U^*)$ consisting of a set of states U , an alphabet of events Σ that trigger transitions in \mathcal{R} via the transition function $\delta : U \times \Sigma \rightarrow U$, an initial state u_{-1} , and a set of goal states U^* . There are no outgoing transitions from any goal state $u^* \in U^*$. Task completion RMs additionally define an output scoring function $\sigma : U \times U \rightarrow \mathbb{R}$, where $\sigma(u, u') = 1$ whenever $u \notin U^*$ and $u' \in U^*$, and 0 otherwise. Task completion RMs are

Mealy machines [20] where reaching the goal state represents a valid completion of the task. We remark that δ and σ are partial functions defined on subsets of $U \times \Sigma$ and $U \times U$. The transition operator for a task completion RM takes in single events e in as input; when multiple events occur simultaneously, the events are passed in sequence to the RM in arbitrary order.

A task completion RM is connected to a Markov game by the labeling function $L^{\mathcal{G}}$. We can project a trajectory in \mathcal{G} to a trajectory in \mathcal{R} by applying $L^{\mathcal{G}}$ to each state, creating a sequence of event sets $(\{e_1 \dots e_k\}_0, \dots, \{e_1 \dots e_k\}_\tau)$. This sequence of event sets transitions \mathcal{R} to create a trajectory (u_0, \dots, u_τ) , where u_0 is the state resulting from $\delta(u_{-1}, \{e_1 \dots e_k\}_0)$, and so forth. We say \mathcal{R} *accepts* a trajectory if $u_\tau \in U^*$. By definition, the cumulative score of σ will be 1 for accepting trajectories, and 0 for all non-accepting trajectories.

2.3 Using task completion RMs in MARL

Recall our problem setting of *centralized training, decentralized execution*: each agent will receive independent observations during execution. In other words, each agent will view the cooperative Markov game as an MDP, where the presence and actions of other agents are captured by the dynamics of their respective environments. To train our agent team, we can use the acceptance condition of a task completion RM as a reward function for MARL. This objective can be naturally compiled down to a reward function expressed over a **product MDP** for individual agent policies $\pi_1 \dots \pi_n$. Concretely, a product MDP synchronizes \mathcal{M} and \mathcal{R} so that an agent may learn a policy over the joint space by coupling the reward machine state u_t during a trajectory with the MDP’s state s_t , producing an action conditioned on both: $\pi(a_t | s_t, u_t)$. If the agents transition to a goal state in \mathcal{R} , they will receive a reward of one; all other transitions will receive a reward of zero. Existing deep RL algorithms, such as PPO [31], have shown success in learning performant policies for RMs in a variety of single-agent settings by learning policies over the product MDP [19, 42].

However, in a MARL setting, providing the full task completion RM for each agent creates difficulties in learning. If all agents see the same states and transitions in \mathcal{R} , then all agents will receive credit when the task is completed, even if one or more agents did not contribute to completion of the task. Recall our running example in Figure 2. Suppose Agent 1 and Agent 2 visit HQ at the same time, and then Agent 2 visits the yellow and red stations, while Agent 3 remains stationary. Because the task is completed, all agents will receive equal reward for that episode. As a result, Agent 3 will think that its stationary behavior is desirable and be encouraged to act similarly in future episodes.

Existing work addresses this shortcoming by introducing the notion of *task decomposition* for a given task completion RM [21]. A decomposition of a task completion RM creates sub-tasks in the form of n smaller RMs $\mathcal{R}_1 \dots \mathcal{R}_n$, derived from the original, that can then be assigned to each agent. Each agent is only concerned with accomplishing their own sub-task encoded by the RM assigned to them. In order to compute a decomposition, a practitioner provides *Local Event Sets* (LES) for each agent. An LES is a subset of events $\Sigma_i \in \Sigma$ that is deemed relevant to agent i ’s individual sub-task and restricts agent i to only observing events in Σ^i . A task completion RM is then *projected* onto these local event sets to create

an individual's sub-task reward machine $\mathcal{R}_i = (U_i, u_{-1^i}, \Sigma_i, \delta_i, U_i^*)$. The states, initial state, and goal states of \mathcal{R}_i are sets of equivalence classes over states in \mathcal{R} based on an equivalence relation that subsumes states whose transitions do not contain any event in Σ_i . The transition function is a projection of the original δ where a transition between two states u_j and u_k in \mathcal{R}_i exists only if an event in Σ_i triggered a transition between two distinct states in \mathcal{R} that were subsumed by u_j and u_k respectively. For the exact procedure of this projection, see [21].

For each agent's sub-task RM \mathcal{R}_i , we define an individual labeling function L_i that projects the set of environment events returned by L to the events belonging to an agent's local event set Σ_i . We say an event e is a *shared event* if it belongs to more than one local event set. For example, the event "Signal" in Figure 2 is a shared event. In the case of shared events, agents' sub-task RMs must *synchronize* on this event. This means that a synchronized event must trigger a transition for all agents' RMs that share the event (i.e. all sub-task RMs must be in the appropriate state for the synchronized event to cause a transition), or no transition is taken from encountering that event for any agent's RM.

The decomposition approach outlined in [21] relies on a practitioner manually designing the local event sets to assign to each agent. In order to help guide the search for an appropriate assignment of local event sets, [21] provides a notion of *validity* for a given decomposition: a decomposition is valid if and only if the parallel composition of all reward machines $\mathcal{R}_1 \dots \mathcal{R}_n$ is bisimilar to the original \mathcal{R} . If a decomposition is valid, then for any trajectory of events ξ , \mathcal{R} accepts ξ if and only if all sub-task RMs $\mathcal{R}_1 \dots \mathcal{R}_n$ accept ξ . In other words, a trajectory of events that accomplishes all decomposed sub-tasks encoded by $\mathcal{R}_1, \dots, \mathcal{R}_n$ is guaranteed to solve the overall task.

3 METHODOLOGY

In practice, multiple valid decompositions often exist for a given RM. However, a valid decomposition may not be *feasible* under the dynamics of a given MDP; that is, the resulting learned policies may not be able to achieve the individual tasks prescribed by a decomposition. Moreover, when multiple feasible decompositions do exist, we do not know which decomposition most efficiently achieves the task or provides the best credit assignment for learning. Recall the Repairs environment and task in Figure 2. A feasible decomposition of the task would be for Agent 1 and Agent 2 to meet at HQ, then for Agent 2 to visit the yellow, then red stations in that order. However, this decomposition is less efficient than a decomposition where Agents 1 and 3 visit HQ, then Agent 2 visits the yellow station and leaves the hazardous region while Agent 3 visits the red station (assuming Agents complete their sub-tasks with optimal efficiency).

Recent work attempts to automate the search for RM decompositions by leveraging additional information provided by a practitioner [34]. In this work, subsets of events in Σ can be provided that either require or forbid an event to belong to a specific agent's local event set Σ_i , along with a utility function that quantifies how valuable an event would be to a specific agent. This information is then used to find a subset of events in Σ that still leads to a valid decomposition of \mathcal{R} , if one exists. However, the aforementioned

approach does not leverage knowledge gained about the MDP during training and therefore cannot ensure that the decomposition generated from their method is feasible or optimally efficient. In what follows, we will introduce our approach, which aims to find the optimal decomposition by learning a policy on-the-fly for many possible sub-tasks during training.

Our approach can be broken down into three primary components: (1) automatically generating a set of possible (candidate) decompositions for our task, (2) using a task-conditioned policy to generalize learning across multiple decompositions, and (3) employing the Upper Confidence Bounds (UCB) strategy [5] to balance exploration and exploitation of candidates throughout training.

3.1 Generating candidate decompositions

Approaches to procedurally generate decompositions of reward machines and similar automaton-based task specifications have been explored in the literature [18, 34]. Such methods can be used to generate a finite set of candidate decompositions $\mathcal{D} = \{(\mathcal{R}_1^1 \dots \mathcal{R}_n^1), \dots, (\mathcal{R}_1^{|\mathcal{D}|} \dots \mathcal{R}_n^{|\mathcal{D}|})\}$ from which the optimal decomposition can be found. We will denote an individual decomposition in our set as $d \in \mathcal{D}$. In our work, we will use the decomposition approach for task completion RMs introduced in [34]. In this approach, all possible *valid* decompositions are generated given Σ , and each one is assigned a score based on three factors: (1) minimizing the average number of events assigned to each agent's local event set, (2) maximizing the similarity amongst the sizes of all local event sets, and (3) maximizing the average 'utility' of each agent's local event set, based on a practitioner-provided utility function that maps the assignment of an event to an agent to a scalar value. In our work, we assume that no utility function has been provided. We will therefore enumerate the 'top- k ' valid candidate decompositions based on a weighted sum of scores pertaining to factors (1) and (2), where k is a hyperparameter that sets the number of decompositions we will consider, i.e. $k = |\mathcal{D}|$. Our objective is to find the decomposition $d^* \in \mathcal{D}$ that leads to a learned policy which will most efficiently achieve the original task \mathcal{R} .

3.2 Task-conditioned policies

We consider each agent policy in our team as sharing a *task-conditioned* policy. Instead of learning separate policies for individual sub-tasks, we learn a single policy that outputs actions conditioned on a specific sub-task RM belonging to a decomposition generated in section 3.1. We can then deploy separate copies of our policy on each agent during execution time.

Policy Architecture. Our policy is represented by a feedforward neural network that is shared and learned across all agents. The neural policy receives four inputs, depending on the individual agent:

- an observation from the underlying MDP (different for each agent);
- an encoding of the selected decomposition $d_j \in \mathcal{D}$ for the current execution episode (same for each agent);
- the current state in their sub-task $\mathcal{R}_i^j \in d_j$ (different for each agent);




Figure 3: A visualization of the information each agent receives using the policy architecture described in section 3.2 for the Repairs task from Figure 2. In addition to the observation gathered from the MDP, each agent’s policy is conditioned on (1) the current state of the original RM task, (2) which decomposition is currently selected, and (3) the current state of their assigned sub-task RM within the selected decomposition.

- an encoding of the overall task and the current state in the overall task (same for each agent).

We visualize these inputs in Figure 3 for an example decomposition of our Repairs task. Learning a task-conditioned policy allows us to distinguish between different sub-tasks within different decompositions while exploiting the generalization capabilities of multi-task learning [24, 39] as well as the natural curriculum learning inherent in decomposition exploration. Let us return to the running example from Figure 2. Imagine that as part of a decomposition selected during training, agent A_1 is only assigned the task of meeting at HQ and nothing else. Even if this decomposition is not optimal, succeeding in this sub-task teaches A_1 the valuable skill of how to reach HQ. This experience will be useful in the future when A_1 is tasked with more complicated sub-tasks, such as “first go to HQ and then the red station”.

Conditioning on the overall task allows us to relax the assumption of independent dynamics as we will describe in Section 3.4. In our experiments, we use one-hot encodings of the selected sub-task and the current position in the sub-task during the rollout. However, the embeddings could in principle be generated in any way, such as learned embeddings from a graph neural network that encode the structure of an RM [46].

3.3 Selecting decompositions during training

A naive approach to our objective of finding an optimal ordered decomposition would be to learn a task-conditioned policy for each $d \in \mathcal{D}$ independently, and then choose the decomposition that performs the best after a certain amount of training. As the number of candidate decomposition increases, this approach quickly becomes intractable. We seek to improve the efficiency of this process by simultaneously learning both the most efficient decomposition and the corresponding policies that achieve this decomposition.

A key insight of our work is that we can use rewards from previous executions of a sub-task \mathcal{R}_i^j to estimate the satisfaction

likelihood of that sub-task. These estimates, which we call *value estimates*, are used as a heuristic in decomposition selection to assess how well a policy is performing on a specific sub-task. Our approach computes value estimates as an *exponential weighted moving sum* of previous rewards, as more recent rewards are typically a more accurate reflection of the performance of the current policy.

We will denote the value estimates for a sub-task \mathcal{R}_i^j as $V_{\mathcal{R}_i^j}$. On episode H of training on decomposition j , we can compute a sub-task’s value estimate as $V_{\mathcal{R}_i^j} = \sum_{h=0}^H \alpha^{H-h} r_h$. Here, r_h represents the reward achieved by agent i under sub-task \mathcal{R}_i from the h -th execution and α is a hyperparameter defining the decay rate.

With value estimates in hand for each agent policy in our team, on each sub-task, we can consider a number of heuristic approaches to utilize them. In this work, we use the Upper Confidence Bound (UCB) algorithm [5], which balances exploring different decompositions with exploiting higher scoring decomposition via a hyperparameter β . The score assigned to each decomposition d_j is an average of the decomposition’s current value estimates for each sub-task $\{V_{\mathcal{R}_1^j}, \dots, V_{\mathcal{R}_n^j}\}$.

We note that the value estimates early in training may be arbitrarily inaccurate due to the lack of progress made in learning the policy for different sub-tasks by each agent. When inaccuracy is high early in training, exploration is critical so that agents can sufficiently optimize towards achieving each candidate sub-task before the value estimates are exploited to converge on the optimal decomposition. Once the optimal decomposition is converged upon, additional training will further optimize the policy conditioned on its corresponding sub-tasks.

3.4 Dealing with dependent dynamics

Prior work in RM-guided MARL [21, 34] assumes that the underlying dynamics of the MDP are independent between agents. This prevents the use of an MDP with dynamics that model collisions

or interactions between agents unless the interaction is explicitly modeled by the task’s reward machine, which requires knowing about the interaction a priori.

Consider the motivating example visualized in Figure 2. The example includes dependent dynamics between the agents in the form of the hazardous region. Only one agent can occupy the region at a time, meaning that an agent’s ability to enter the region is dependent on the other agents’ positions and actions.

Prior works [21, 34] require the assumption of independent dynamics so that the completion of an individual agent’s reward machine is independent of the behavior of other agents. With dependent dynamics, this is rarely the case. For example, imagine that Agent 3 is assigned to reach the yellow target and Agent 2 is assigned to reach the red target. If Agent 3 enters the hazardous region and reaches the yellow target, it accomplishes its sub-task. Now, completion of the overall task only requires Agent 2 to reach the red target. Since the red target is located in the hazardous region, this requires Agent 3 to exit the hazardous region. However, Agent 3 has already accomplished its sub-task (and has no knowledge of Agent 2’s task), so it has no incentive to leave the hazardous region.

We are able to relax the assumption of independent dynamics made in previous work [21, 34] and handle environments such as the one in Figure 2 by allowing agents to condition not only on the status of their sub-task, but on the status of the overall task as well. In addition, agents are given a small reward for the completion of the overall task during training along with their primary reward for completing their assigned sub-task. This reward structure incentivizes agents to condition on the status of other agent’s tasks (e.g., that Agent 2 needs to go to red) so that the overall task can be completed (e.g., Agent 3 is incentivized to leave the hazardous region).

4 EXPERIMENTS

We evaluate our proposed approach, which we refer to as **LOTaD** (Learning Optimal Task Decompositions), in a variety of MARL settings with varying task complexity. Our code is available at <https://github.com/thomasychen/LOTaD>.

In our experiments, we seek to answer the following research questions: **(RQ1)** Does LOTaD outperform existing approaches to RM-guided MARL that are not learning-informed, including approaches that do not perform task decomposition? **(RQ2)** Can LOTaD learn to avoid pitfalls in learning that may occur due to dependent dynamics across agents? **(RQ3)** How does varying the number of candidate decompositions $k = |\mathcal{D}|$ affect the learning performance of LOTaD?

4.1 Environments and tasks

Our environments include the **Repairs** environment and RM task presented in Figure 2. The environment is instantiated as a 7x12 grid in which agents can move in any of the cardinal directions, or take a no-op action where no movement occurs. We make the Repairs environment stochastic by giving a small chance of an agent “slipping” and taking a random action rather than the action selected by their policy. Each agent observes only their position in the world at each timestep.

In addition to the Repairs environment, we also include two “Buttons” environments that require a team of agents to press a series of buttons in a particular order. We instantiate two environment-task pairs: First, we use **Cooperative Buttons**, a task from [21] where any agent must reach a specified goal location, but in order to do so, must traverse regions that can only be crossed once a corresponding button has been pressed. Second, we use **Four-Buttons**, a task where two agents must press four buttons (yellow, green, blue, and red) in an environment, with an ordering constraint that the yellow button must be pressed before the red button. Both of these environments are represented as 10x10 grids with the same observation space, action space, and “slipping” dynamics as the Repair environment. Visualizations of these environments are presented in our Appendix [32].

Lastly, we evaluate LOTaD on two environments from the popular multi-agent benchmark *Overcooked* [9]. In both environments, we provide the same simple task of delivering a soup to the delivery station, which requires putting three onions in a pot, plating, and delivering the soup. Agents must coordinate depending on the dynamics of their environment to efficiently cook and deliver the soup. We use a **Cramped-Corridor** environment, where two agents must navigate around one another to reach the pot at the end of a small corridor in a cramped room, and an **Asymmetric-Advantages** environment, where two agents are in separate rooms, but have access to an asymmetric set of resources in their respective rooms. We visualize both Overcooked environments in our Appendix [32]. The observation space and action space for these environments are the same as those provided in the Overcooked implementation from [27] with the adjustment that we do not expose the number of onions in the pot or whether the soup has finished cooking in order to ensure that no information provided by the labeling function is redundant in the agents’ observations.

In all environments, we apply a discount factor $\gamma < 1.0$ to the reward offered by our RM to incentivize our agents to accomplish the task as efficiently as possible. We generate $k = 10$ decomposition candidates in all experiments using the generation method described in Section 3.1 for LOTaD to search amongst. We provide additional information regarding each environment and task in our Appendix [32].

4.2 Baselines

To evaluate LOTaD, we compare against a baseline that selects a decomposition using the **ATAD** method [34]. This approach selects a set decomposition prior to learning based on the scoring method described in Section 3.1. We break ties between top scoring decompositions arbitrarily. In addition to this baseline, we also compare against a baseline approach that assigns each agent the overall task. We call this baseline **Monolithic** and use it to evaluate how MARL would perform if no decomposition of the RM was used. We use PPO [31] with a Gaussian policy over the action space for each agent as our RL algorithm in every environment.

4.3 Results

We plot training curves for all experimental domains in Figure 4. In these curves, we report the current best discounted reward achieved

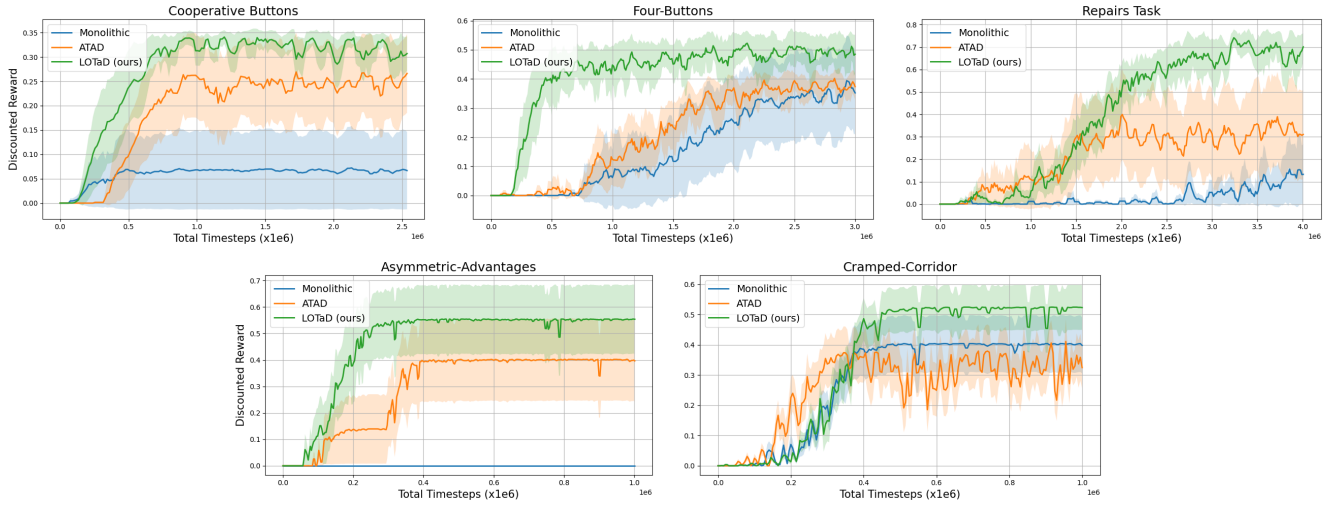


Figure 4: Training curves for LOTaD and baseline methods in our experimental domains. Results are averaged over 5 random seeds.




Figure 5: Training curves for LOTaD in the Repairs Task and Cramped-Corridor environments demonstrating the effect of conditioning on the overall task state along with individual sub-task states for each agent.

by LOTaD amongst all decomposition candidates under consideration. We find that LOTaD outperforms both the monolithic and ATAD baselines across all environments, answering (RQ1) in the affirmative. In many of our environments, learning decompositions with LOTaD allows for agents to explicitly parallelize their contributions towards achieving the task. For example, in the Four-Buttons environment, a candidate decomposition allows for one agent to visit the yellow and green buttons while the other agent visits the blue button.

Interestingly, LOTaD-learned decompositions are useful even when explicit parallelization is not possible, such as the Overcooked RM task. Decompositions of this task can still facilitate multi-agent learning: for example, one agent may be tasked with putting all three onions in the pot, while the other agent must wait for the soup to be cooked before plating and delivering the soup. In this decomposition, the agent tasked with plating and delivering the

soup may fetch a plate and stand near the pot so that the soup can be quickly delivered upon completion of cooking.

In comparison to LOTaD, the ATAD baseline is unable to consistently find performant task decompositions. Although ATAD selects a decomposition based on the same scoring method by which we select \mathcal{D} , there exist many possible decompositions tied for the highest achievable score in a given task. As a result, ATAD may select the optimal decomposition when the optimal decomposition is also the highest scoring, but performs suboptimally when this is not the case. In addition to lower reward, this leads to a higher variance in performance by ATAD, as evidenced by Figure 4. The Monolithic baseline consistently achieves the lowest reward due to the sparsity of the overall task reward. Similar to ATAD, we notice that policy learning with the Monolithic baseline is unreliable and tends to converge more slowly than LOTaD.

We find that LOTaD is able to successfully accomplish the task in the Repairs and Overcooked Cramped-Corridor environments,




Figure 6: Training curves for the UCB selection strategy on increasingly sized \mathcal{D} for the Four-Buttons task. Results are averaged over 10 random seeds.

which both involve dependent dynamics amongst agents. LOTaD is indeed able to avoid issues when training agent teams in environments with dependent dynamics, affirming (RQ2). To further investigate our hypothesis of whether a global task view alleviates these issues, we ran LOTaD *without* an encoding of the overall task in the Repairs and Cramped-Corridor environment. We present the training curves from this ablation study in Figure 5. We find that without the overall task encoding, LOTaD is largely unable to accomplish the Repairs task, and is slower and less stable in accomplishing the Cramped-Corridor task. This suggests that the overall task encoding enables our agents to learn policies that progress towards completion of the overall task even when an agent’s individual sub-task is already achieved. In our other environment settings, where codependent agent dynamics do not exist, removing the overall task encoding does not significantly affect training.

To answer (RQ3), we perform ablation studies where we vary the size of \mathcal{D} . These training curves are visualized in Figures 6. We see that values of k that are either very small or very large make the learning problem more challenging for LOTaD. We reason that if k is too large, more exploration is required, and LOTaD may struggle to find the optimal decomposition amongst many candidates. If k is too small, there is a lower chance of “good” decompositions appearing in the set \mathcal{D} , which inherently limits the performance potential of LOTaD. We note that the inherent randomness of LOTaD, as well as the choice of environment, task, and the amount of exploration prioritized by our UCB hyperparameter β , may confound results.

5 RELATED WORK

Automaton-based task specifications in RL. An extensive body of work has explored automaton-based task specification for RL agents. Previous efforts have proposed RL approaches to policy learning for reward machines [8, 14, 38] or automata-based representations of temporal logic [3, 12, 13, 15, 28, 33, 41, 42] by augmenting the state space of the MDP. These efforts focus primarily on single-agent settings and are not designed to handle the MARL case with shared objectives or formally decompose the automaton.

(Symbolic) task decomposition for multi-agent teams. Symbolic structures have been leveraged to facilitate efficient multi-agent learning in a variety of settings. A number of these approaches

rely on a known dynamics model [16, 29, 30] for planning-based approaches. In contrast, we assume no prior knowledge of the environment dynamics as is standard in reinforcement learning.

Closely related to our contributions, previous works explored learning decompositions of tasks for teams of agents. Some previous works explored decomposing symbolic tasks by hand designing decompositions or using task- and environment-agnostic heuristics [21, 34]. However, these works either require extensive human involvement in determining the optimal decomposition or do not offer a way to choose a decomposition based on MDP dynamics. Moreover, these approaches are limited to MDPs in which the dynamics of agents are independent. Other works learn role assignments for traditional reward functions (i.e., non-symbolic) with multi-agent teams [44] or decompose traditional (non-Markovian) reward functions [35] for teams of agents but these methods do not easily extend to non-Markovian rewards which we consider.

Credit assignment in multi-agent RL. Credit assignment, originally explored in the single-agent setting [11, 26, 37, 40], is a problem that has been extensively explored in multi-agent learning literature and can be divided between *explicit* and *implicit* solutions. Explicit credit assignment most closely resembles our work, typically assuming that value functions are given in a specific form that allows certain types of decompositions, such as additive decompositions [22], or value factorization [36, 43]. Other explicit methods are based on assuming a hierarchical execution structure [2, 10, 25]. Implicit credit assignment instead attempts to perform credit assignment without an explicit structure, for example, by deriving individual policy gradients for each agent derived from a centralized critic [47].

6 CONCLUSION

We introduce a novel approach for learning the optimal decomposition of a task completion RM for MARL through model free interactions with the environment. Our framework simultaneously learns the optimal decomposition and the policies that solve that decomposition amongst a set of candidates. Our experiments show that we improve the sample efficiency of multi-agent learning in model-free deep RL settings even when the optimal decomposition is not known a priori. Through our ablations, we have shown that both including the encoding of the overall task and intelligent decomposition selection is critical for sample efficient learning of reward machines in multiagent settings.

We see a number of compelling directions for future work. For example, we are interested in exploring how different representations of an RM may enable greater sample efficiency by exploiting semantic similarity amongst overlapping sub-tasks. Lastly, we are interested in exploiting the inherent curriculum present in decomposition exploration by generalizing our multi-agent setting to solving multiple tasks, building on prior work in goal-conditioned RL on automata-based tasks [24, 39].

ACKNOWLEDGMENTS

This work was supported in part by DARPA Contract FA8750-23-C-0080 (ANSR) and the DARPA Contract HR00112490425 (TIAMAT), by Nissan and Toyota under the iCyPhy center, and an NSF Graduate Fellowship.

REFERENCES

- [1] Adrian K. Agogino and Kagan Tumer. 2004. Unifying Temporal and Structural Credit Assignment Problems. In *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, 19–23 August 2004, New York, NY, USA. IEEE Computer Society, 980–987. <https://doi.org/10.1109/AAMAS.2004.10098>
- [2] Adrian K Agogino and Kagan Tumer. 2004. Unifying temporal and structural credit assignment problems. In *Autonomous agents and multi-agent systems conference*.
- [3] Rajeev Alur, Suguman Bansal, Osbert Bastani, and Kishor Jothimurugan. 2022. A Framework for Transforming Specifications in Reinforcement Learning. In *Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday (Lecture Notes in Computer Science, Vol. 13660)*, Jean-François Raskin, Krishnendu Chatterjee, Laurent Doyen, and Rupak Majumdar (Eds.), Springer, 604–624. https://doi.org/10.1007/978-3-031-22337-2_29
- [4] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. 2009. Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science* 410, 19 (2009), 1876–1902.
- [5] Peter Auer. 2002. Using Confidence Bounds for Exploitation-Exploration Tradeoffs. *J. Mach. Learn. Res.* 3 (2002), 397–422. <http://jmlr.org/papers/v3/auer02a.html>
- [6] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47 (2002), 235–256.
- [7] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research* 27, 4 (2002), 819–840.
- [8] Alberto Camacho, Rodrigo Toro Icarte, Torny Q. Klassen, Richard Valenzano, and Sheila A. McIlraith. 2019. LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 6065–6073. <https://doi.org/10.24963/ijcai.2019/840>
- [9] Micah Carroll, Rohin Shah, Mark K. Ho, Tom Griffiths, Sanjit A. Seshia, Pieter Abbeel, and Anca D. Dragan. 2019. On the Utility of Learning about Humans for Human-AI Coordination. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.), 5175–5186. <https://proceedings.neurips.cc/paper/2019/hash/f5b1b89d98b7286673128a5fb112cb9a-Abstract.html>
- [10] Lei Feng, Yuxuan Xie, Bing Liu, and Shuyan Wang. 2022. Multi-level credit assignment for cooperative multi-agent reinforcement learning. *Applied Sciences* 12, 14 (2022), 6938.
- [11] Johan Ferret, Raphaël Marinier, Matthieu Geist, and Olivier Pietquin. 2019. Self-attentional credit assignment for transfer in reinforcement learning. *arXiv preprint arXiv:1907.08027* (2019).
- [12] Jie Fu and Ufuk Topcu. 2014. Probably Approximately Correct MDP Learning and Control With Temporal Logic Constraints. In *Robotics: Science and Systems X, University of California, Berkeley, USA, July 12–16, 2014*, Dieter Fox, Lydia E. Kavraki, and Hanna Kurniawati (Eds.), <https://doi.org/10.15607/RSS.2014.X.039>
- [13] Mohammadhosein Hasanbeig, Daniel Kroening, and Alessandro Abate. 2020. Deep reinforcement learning with temporal logics. In *International Conference on Formal Modeling and Analysis of Timed Systems*. Springer, 1–22.
- [14] Rodrigo Toro Icarte, Torny Q Klassen, Richard Valenzano, and Sheila A McIlraith. 2022. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research* 73 (2022), 173–208.
- [15] Kishor Jothimurugan, Rajeev Alur, and Osbert Bastani. 2019. A composable specification language for reinforcement learning tasks. *Advances in Neural Information Processing Systems* 32 (2019).
- [16] Mohammad Karimadini and Hai Lin. 2011. Cooperative Tasking for Deterministic Specification Automata. *Asian Journal of Control* 18 (01 2011). <https://doi.org/10.1002/asjc.1300>
- [17] Michael N Katehakis and Arthur F Veinott Jr. 1987. The multi-armed bandit problem: decomposition and computation. *Mathematics of Operations Research* 12, 2 (1987), 262–268.
- [18] Niklas Lauffer, Beyazıt Yalcinkaya, Marcell Vazquez-Chanlatte, Ameesh Shah, and Sanjit A. Seshia. 2022. Learning Deterministic Finite Automata Decompositions from Examples and Demonstrations. In *22nd Formal Methods in Computer-Aided Design, FMCAD 2022, Trento, Italy, October 17–21, 2022*, Alberto Griggio and Neha Rungta (Eds.), IEEE, 1–6. https://doi.org/10.34727/2022/ISBN.978-3-85448-053-2_39
- [19] Andrew C. Li, Zizhao Chen, Torny Q. Klassen, Pashootan Vaezipoor, Rodrigo Toro Icarte, and Sheila A. McIlraith. 2024. Reward Machines for Deep RL in Noisy and Uncertain Environments. *CoRR abs/2406.00120* (2024). <https://doi.org/10.48550/ARXIV.2406.00120> arXiv:2406.00120
- [20] George H Mealy. 1955. A method for synthesizing sequential circuits. *The Bell System Technical Journal* 34, 5 (1955), 1045–1079.
- [21] Cyrus Neary, Zhe Xu, Bo Wu, and Ufuk Topcu. 2021. Reward Machines for Cooperative Multi-Agent Reinforcement Learning. In *AAMAS ’21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3–7, 2021*, Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé (Eds.), ACM, 934–942. <https://doi.org/10.5555/3463952.3464063>
- [22] Duc Thien Nguyen, Akshat Kumar, and Hoang Chuin Lau. 2017. Policy gradient with value function approximation for collective multiagent planning. *Advances in neural information processing systems* 30 (2017).
- [23] Afshin Oroojlooy and Davood Hajinezhad. 2023. A review of cooperative multi-agent deep reinforcement learning. *Appl. Intell.* 53, 11 (2023), 13677–13722. <https://doi.org/10.1007/S10489-022-04105-Y>
- [24] Wenjie Qiu, Wensen Mao, and He Zhu. 2023. Instructing Goal-Conditioned Reinforcement Learning Agents with Temporal Logic Objectives. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10–16, 2023*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.), http://papers.nips.cc/paper_files/paper/2023/hash/7b35a69f434b5eb07ed1b1ef16ace52c-Abstract-Conference.html
- [25] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research* 21, 178 (2020), 1–51.
- [26] Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degraeve, Tom Wiele, Vlad Mnih, Nicolas Heess, and Jost Tobias Springenberg. 2018. Learning by playing solving sparse reward tasks from scratch. In *International conference on machine learning*. PMLR, 4344–4353.
- [27] Alexander Rutherford, Benjamin Ellis, Matteo Gallici, Jonathan Cook, Andrei Lupu, Gardar Ingvarsson, Timon Willi, Akbir Khan, Christian Schroeder de Witt, Alexandra Souly, Saptarashmi Bandyopadhyay, Mikayel Samvelyan, Minqi Jiang, Robert Tjarko Lange, Shimon Whiteson, Bruno Lacerda, Nick Hawes, Tim Rocktaschel, Chris Lu, and Jakob Nicolaus Foerster. 2023. JaxMARL: Multi-Agent RL Environments in JAX. *arXiv preprint arXiv:2311.10090* (2023).
- [28] Dorsa Sadigh, Eric S Kim, Samuel Coogan, S Shankar Sastry, and Sanjit A Seshia. 2014. A learning based approach to control synthesis of markov decision processes for linear temporal logic specifications. In *53rd IEEE Conference on Decision and Control*. IEEE, 1091–1096.
- [29] Philipp Schillinger, Mathias Bürger, and Dimos V. Dimarogonas. 2016. Decomposition of Finite LTL Specifications for Efficient Multi-agent Planning. In *Distributed Autonomous Robotic Systems, The 13th International Symposium, DARS 2016, Natural History Museum, London, UK, November 7–9, 2016 (Springer Proceedings in Advanced Robotics, Vol. 6)*, Roderich Groß, Andreas Kolling, Spring Berman, Emilio Frazzoli, Alcherio Martinoli, Fumitoshi Matsuno, and Melvin Gauci (Eds.), Springer, 253–267. https://doi.org/10.1007/978-3-319-73008-0_18
- [30] Philipp Schillinger, Mathias Bürger, and Dimos V. Dimarogonas. 2018. Simultaneous task allocation and planning for temporal logic goals in heterogeneous multi-robot systems. *Int. J. Robotics Res.* 37, 7 (2018), 818–838. <https://doi.org/10.1177/0278364918774135>
- [31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR abs/1707.06347* (2017). arXiv:1707.06347 <https://arxiv.org/abs/1707.06347>
- [32] Ameesh Shah, Niklas Lauffer, Thomas Chen, Nikhil Pitta, and Sanjit A. Seshia. 2025. Learning Symbolic Task Decompositions for Multi-Agent Teams. arXiv:2502.13376 [cs.LG] <https://arxiv.org/abs/2502.13376>
- [33] Ameesh Shah, Cameron Voloshin, Chenxi Yang, Abhinav Verma, Swarat Chaudhuri, and Sanjit A. Seshia. 2024. Deep Policy Optimization with Temporal Logic Constraints. *CoRR abs/2404.11578* (2024). <https://doi.org/10.48550/ARXIV.2404.11578> arXiv:2404.11578
- [34] Sophia Smith, Cyrus Neary, and Ufuk Topcu. 2023. Automatic Decomposition of Reward Machines for Decentralized Multiagent Reinforcement Learning. In *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 5423–5430.
- [35] Changyin Sun, Wenzhang Liu, and Lu Dong. 2020. Reinforcement learning with task decomposition for cooperative multiagent systems. *IEEE transactions on neural networks and learning systems* 32, 5 (2020), 2054–2065.
- [36] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 2017. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296* (2017).
- [37] Richard Stuart Sutton. 1984. *Temporal credit assignment in reinforcement learning*. University of Massachusetts Amherst.
- [38] Rodrigo Toro Icarte, Torny Q. Klassen, Richard Valenzano, and Sheila A. McIlraith. 2022. Reward Machines: Exploiting Reward Function Structure in Reinforcement Learning. *J. Artif. Int. Res.* 73 (may 2022), 36. <https://doi.org/10.1613/jair.1.12440>
- [39] Pashootan Vaezipoor, Andrew C Li, Rodrigo A Toro Icarte, and Sheila A McIlraith. 2021. Ltl2action: Generalizing ltl instructions for multi-task rl. In *International Conference on Machine Learning*. PMLR, 10497–10508.
- [40] Hado van Hasselt, Sephora Madjiheurem, Matteo Hessel, David Silver, André Barreto, and Diana Borsa. 2021. Expected eligibility traces. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 9997–10005.

- [41] Cameron Voloshin, Hoang Minh Le, Swarat Chaudhuri, and Yisong Yue. 2022. Policy Optimization with Linear Temporal Logic Constraints. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). <https://openreview.net/forum?id=yZcPRIZEwOG>
- [42] Cameron Voloshin, Abhinav Verma, and Yisong Yue. 2023. Eventual Discounting Temporal Logic Counterfactual Experience Replay. *arXiv:2303.02135 [cs.LG]*
- [43] Jianhao Wang, Zhizhou Ren, Beining Han, Jianing Ye, and Chongjie Zhang. 2021. Towards understanding cooperative multi-agent q-learning with value factorization. *Advances in Neural Information Processing Systems* 34 (2021), 29142–29155.
- [44] Tonghan Wang, Tarun Gupta, Anuj Mahajan, Bei Peng, Shimon Whiteson, and Chongjie Zhang. 2020. Rode: Learning roles to decompose multi-agent tasks. *arXiv preprint arXiv:2010.01523* (2020).
- [45] Zhe Xu, Ivan Gavran, Yousef Ahmad, Rupak Majumdar, Daniel Neider, Ufuk Topcu, and Bo Wu. 2020. Joint Inference of Reward Machines and Policies for Reinforcement Learning. In *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling, Nancy, France, October 26-30, 2020*, J. Christopher Beck, Olivier Buffet, Jörg Hoffmann, Erez Karpas, and Shirin Sohrabi (Eds.). AAAI Press, 590–598. <https://ojs.aaai.org/index.php/ICAPS/article/view/6756>
- [46] Beyazit Yalcinkaya, Niklas Lauffer, Marcell Vazquez-Chanlatte, and Sanjit Seshia. [n.d.]. Automata Conditioned Reinforcement Learning with Experience Replay. In *NeurIPS 2023 Workshop on Goal-Conditioned Reinforcement Learning*.
- [47] Meng Zhou, Ziyu Liu, Pengwei Sui, Yixuan Li, and Yuk Ying Chung. 2020. Learning implicit credit assignment for cooperative multi-agent reinforcement learning. *Advances in neural information processing systems* 33 (2020), 11853–11864.