

Logical Specifications-guided Dynamic Task Sampling for Reinforcement Learning Agents

Yash Shukla¹, Tanushree Burman¹, Abhishek N. Kulkarni²,
Robert Wright³, Alvaro Velasquez⁴, Jivko Sinapov¹

¹Tufts University

²University of Texas at Austin

³Georgia Tech Research Institute

⁴University of Colorado Boulder

Abstract

Reinforcement Learning (RL) has made significant strides in enabling artificial agents to learn diverse behaviors. However, learning an effective policy often requires a large number of environment interactions. To mitigate sample complexity issues, recent approaches have used high-level task specifications, such as Linear Temporal Logic (LTL_f) formulas or Reward Machines (RM), to guide the learning progress of the agent. In this work, we propose a novel approach, called Logical Specifications-guided Dynamic Task Sampling (LSTS), that learns a set of RL policies to guide an agent from an initial state to a goal state based on a high-level task specification, while minimizing the number of environmental interactions. Unlike previous work, LSTS does not assume information about the environment dynamics or the Reward Machine, and dynamically samples promising tasks that lead to successful goal policies. We evaluate LSTS on a grid-world and show that it achieves improved time-to-threshold performance on complex sequential decision-making problems compared to state-of-the-art RM and Automaton-guided RL baselines, such as Q-Learning for Reward Machines and Compositional RL from logical Specifications (DIRL). Moreover, we demonstrate that our method outperforms RM and Automaton-guided RL baselines in terms of sample-efficiency, both in a partially observable robotic task and in a continuous control robotic manipulation task.

1 Introduction

Agents are now capable of learning optimal control behavior for a broad spectrum of tasks, ranging from Atari games to robotic manipulation tasks, thanks to recent advancements in Reinforcement Learning (RL). Despite the progress made in RL, learning an optimal task policy using model-free RL techniques still suffers from sample complexity issues because of sparse reward settings and unknown transition dynamics (Lattimore, Hutter, and Sunehag 2013). These challenges further intensify in long-horizon settings, where the agent needs to perform a series of correct sequential decisions to achieve the goal. Certain tasks (such as - robot needs to make dinner only if it bought groceries in the afternoon) require the agent to *encode* and *remember* its episodic history (whether the groceries were bought) in order to solve the task effectively. To alleviate this issue in complicated

tasks, several lines of work have explored representing the goal using an intricately shaped reward function that guides the agent toward the goal (Grzes 2017). However, generating such a reward function requires the human engineer to assign ‘importance’ weights to certain aspects of the task, which is time consuming and assumes knowledge on which aspects of the task are important. Poorly engineered reward functions can lead to local optima, where the agent learns to satisfy only a subset of goals and ignores the rest.

Recent research has investigated representing the goal using high-level specification languages, such as finite-trace Linear Temporal Logic (LTL_f) (De Giacomo and Vardi 2013), Reward Machines (RM) (Icarte et al. 2022), SPECTRL (Jothimurugan, Alur, and Bastani 2019) that allow defining the goal of the task using a graphical representation of sub-tasks. The high-level objective is known before commencing the task, and the graphical representation allows the agent to learn policies that achieve easier sub-goals initially, and build upon them to achieve complex goals. Encoding the task using a graphical structure allows us to tackle the problem in a Markovian manner by tracking the history as a part of the state space (Afzal et al. 2023), thereby allowing the agent to keep track of its episodic history. For instance, if the task for a robot is to reach kitchen and then make dinner, the graphical structure of the task obtained from the high-level specification allows the agent to reason whether it has reached the kitchen before it can commence its policy for making dinner. RM approaches still require human guidance in defining the reward structure, which is dependent on knowing how much reward should be assigned for accomplishing each sub-goal. The process of designing the reward structure assumes that the human engineer is aware of how much should reward should the agent receive when it accomplishes the sub-goals in particular order. This assumption is infeasible in scenarios when the structure of the environment or the exact order in which the sub-goals must be achieved is unknown in advance. In contrast, our method does not require access to the reward structure.

Another method, Compositional RL from Logical Specifications (DIRL) (Jothimurugan et al. 2021) mitigates the reward assignment issue by using Dijkstra’s algorithm to determine which sub-tasks (edges) should be explored in the SPECTRL DAG graph (Jothimurugan, Alur, and Bastani 2019) in order to learn policies to reach nodes in the DAG

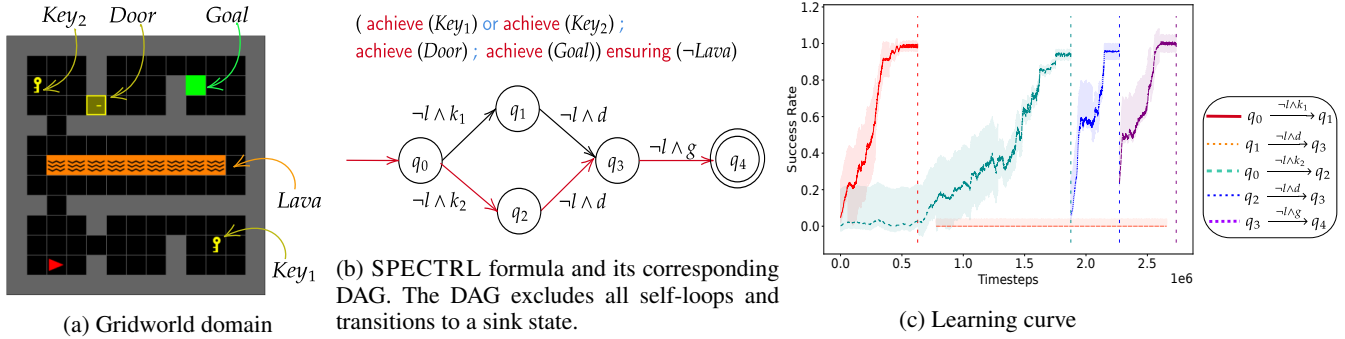


Figure 1: (a) Task environment. The agent (denoted using red triangle) has to *pick-up* one of the keys and *toggle* the door open to reach the goal state; (b) The SPECTRL formula for the task and its DAG. Formulas l , k_1 , k_2 , d and g correspond to *Lava*, *Key₁*, *Key₂*, *Door* and *Goal* respectively; (c) Learning curves for individual sub-tasks (averaged over 10 trials) generated using *LSTS*. The path chosen by *LSTS* is highlighted in red in Fig.1(b)

that yield the highest success rate. DiRL requires the agent to learn RL policies for satisfying *all* outgoing edge propositions (each edge encodes a sub-task) from such nodes. However, this approach requires the agent to explore a sub-task for a manually specified number of interactions, which requires knowledge about the task complexity. DiRL can be sample-inefficient as it attempts to learn unproductive policies on unpromising some sub-tasks, as the agent has to spend the defined number of interactions learning a policy for the sub-task. Unlike DiRL, our approach is sample-efficient as it finds unpromising sub-tasks based on the learning progress of the sub-tasks, and discards them; saving costly interactions and converging to a successful policy faster. This problem of minimizing the overall number of interactions while learning a set of successful policies is non-trivial as the problem equates to finding the shortest path in a graph where the true edge weights are unknown *a priori* (Szepesvári 2004). In our case, the edge weight denotes the total number of environmental interactions required by the agent to learn a successful policy for the sub-task encoded by the edge, in which the agent must induce a visit to a state where certain properties hold true. And, we can sample interactions for a sub-task only if we have a policy to reach the edge’s source node from the start node of the graph, making the learning process more sample-inefficient.

To address the above challenges, we present Logical Specifications-guided Dynamic Task Sampling (*LSTS*). We begin with a high-level objective represented using SPECTRL specification formulas which can equivalently be represented using directed acyclic graphs (DAG) (Jothimurugan, Alur, and Bastani 2019). The DAG structure encodes memory, helping the agent understand what events of interest have occurred in the past, and which events must occur to reach the accepting states. Our key insight is to learn RL policies for sub-tasks defined using cost-to-goal formulation of the edges of the DAG. Specifically, the agent transitions from the node q to p in the DAG when the propositional logic formula labeling the edge (q, p) evaluates to true. We use the set of propositional logic formulas labeling the outgoing edges from a given node in DAG to de-

fine sub-tasks. The trajectory induced by a successful RL policy for the sub-task enables the agent’s high-level state in the DAG to transition from the source node to the destination node of the edge defining the sub-task. We utilize an adaptive Teacher-Student learning approach, where, (1) the Teacher agent uses its high-level policy to actively sample a sub-task for the Student agent to explore. The high-level policy considers the DAG representation and the Student agent’s expected performance on all the sub-tasks. The goal of the Teacher agent is to satisfy the high-level objective in the fewest number of environmental interactions, and (2) The Student agent explores the environment for few interactions (much fewer than the interactions required to learn a successful policy for the sub-task) while simultaneously updating its low-level RL policy for the sampled sub-task. Next, the Teacher observes the Student’s performance and updates its high-level policy. Steps (1) and (2) continue alternately until the Student agent learns a set of successful policies to satisfy the high-level objective.

Running example: As an example, let us look at the environment shown in Fig. 1a. The goal for the agent is to collect *any* of the two *Keys*, followed by opening the *Door* and then reaching the *Goal* while avoiding the *Lava* at all times. The task’s high-level objective (ϕ) is represented using the SPECTRL formula and its corresponding DAG representation \mathcal{G}_ϕ in Fig. 1b. The DAG does not contain information about the environment configuration, such as: the optimal number of interactions required to reach *Door* from *Key₁* are much higher compared to the interactions required to reach *Door* from *Key₂*, making the *Key₁* to *Door* trajectory sub-optimal. A naive approach to tackle this problem would involve learning individual RL policies for all the transitions of DAG, and then later identifying which path of the DAG leads to goal satisfaction in the fewest interactions. However, this approach will still explore all the paths of the DAG and prove sample-inefficient. It is crucial to prevent any additional interactions the agent spends in learning a policy for the sub-task defined by the edge $q_1 \xrightarrow{\neg l \wedge d} q_3$ as the path $q_0 \xrightarrow{\neg l \wedge k_1} q_1 \xrightarrow{\neg l \wedge d} q_3 \xrightarrow{\neg l \wedge g} q_4$ will always be sub-

optimal. In our proposed approach *LSTS*, the Student agent begins with the *aim* of learning two distinct RL policies: π_1 for the task of visiting *Key*₁ and π_2 for the task of visiting *Key*₂, both avoiding *Lava*. The Teacher agent initially samples evenly from these two sub-tasks for the Student but later biases its sampling toward the sub-task on which the Student agent shows higher learning potential. Once the Student agent learns a successful policy for one of the sub-tasks (let’s say the learned policy π_1^* corresponding to the sub-task defined by the transition $q_0 \xrightarrow{\neg l \wedge k_1} q_1$), the Teacher does not sample that task anymore, identifies the next task(s) for the Student using the DAG representation, and appends them to the set of tasks it is currently sampling (in this case, the only next task is: $q_1 \xrightarrow{\neg l \wedge d} q_3$). Since the Student agent only has access to the state distribution over q_0 , it follows the trajectory given by π_1^* to reach a state that lies in the set of states where the proposition $\neg \text{Lava} \wedge \text{Key}_1$ holds true before commencing its learning for the policy (π_3) for $q_1 \xrightarrow{\neg l \wedge d} q_3$. If the Student agent learns the policies π_2^* for satisfying the sub-task defined by $q_0 \xrightarrow{\neg l \wedge k_2} q_2$ and π_4^* for $q_2 \xrightarrow{\neg l \wedge d} q_3$ before learning π_3 , it effectively has a set of policies to reach the node q_3 . Thus, the Teacher will now only sample the next task for the Student in the DAG representation $q_3 \xrightarrow{\neg l \wedge g} q_4$, as learning RL policies for paths that reach q_3 are effectively redundant. This process continues iteratively until the Student agent learns a set of policies that reach the goal node (q_4) from the start node (q_0). The learning curves in Fig. 1c empirically validate the running example. As evident from the learning curves, the Student agent learns policies for the path $q_0 \xrightarrow{\neg l \wedge k_2} q_2 \xrightarrow{\neg l \wedge d} q_3 \xrightarrow{\neg l \wedge g} q_4$ that produce trajectories to reach the goal node q_4 from the initial node q_0 , without excessively wasting interactions on the unpromising sub-task $q_1 \xrightarrow{\neg l \wedge d} q_3$. The dashed lines in Fig. 1c signify the interactions at which a task policy converged.

The dynamic task sampling strategy promotes *LSTS* to achieve sample-efficient learning on complex tasks by identifying unpromising tasks and discarding them, saving costly interactions. Our empirical results show that *LSTS* reduces environmental interactions by orders of magnitude compared to Specifications-Guided RL Baseline D_{IRL}, Reward Machine-based baselines QRM (Icarte et al. 2018), GRS (Camacho et al. 2018), and curriculum learning baseline TSCL (Matiisen et al. 2020). We also evaluate *LSTS*^{ct}, a modified algorithm that further improves sample efficiency by continuing exploration on a new sub-task once the sub-task goal state is reached. We perform evaluation on two simulated robotic tasks and demonstrate that *LSTS* reduces the number of interactions by orders-of-magnitude when compared to automaton-guided RL baselines.

2 Related Work

Automaton-guided RL approaches utilize temporal logic-based language specifications to define tasks (Toro Icarte et al. 2018; Alur et al. 2022; Xu and Topcu 2019; Neary et al. 2021). Separating policies for task sub-goals aids in abstracting knowledge that can be utilized in novel

tasks (Icarte et al. 2018), without reliance on a dense reward function. Another technique is to shape the reward in proportion to the distance from the accepting node in the automaton (Camacho et al. 2018); however, this often leads to suboptimal reward settings. Augmenting the reward function with Monte Carlo Tree Search helps mitigate this issue (Velasquez et al. 2021). This approach requires the ability to plan-ahead in the environment, which is not always feasible. Automaton-guided RL has been used to aid navigational exploration for robotic domains (Cai et al. 2023) and for multi-agent settings (Hammond et al. 2021). Generating a curriculum given the high-level objective (Shukla et al. 2023) requires access to the Object-Oriented MDP (Diuk, Cohen, and Littman 2008), which cannot be obtained if environment details are not known in advance. D_{IRL} interleaves high-level planning with RL to learn a policy for each edge, which overcomes challenges arising from poor representations (Jothimurugan et al. 2021). This approach becomes inefficient in terms of number of interactions, as it requires the agent to act for a predetermined number of interactions, even if learning the task does not show any promise. Unlike previous works, in this paper, we propose an logical specifications-guided dynamic task sampling approach that does not require access to the environment dynamics or the Reward Machine, and efficiently samples tasks that show promise toward the high-level objective, saving interactions on unpromising tasks.

Teacher-Student algorithms (Matiisen et al. 2020) have been previously studied in Curriculum Learning literature (Narvekar et al. 2020; Shukla et al. 2022, 2024) and in the Intrinsic Motivation literature (Oudeyer and Kaplan 2009). These approaches utilize a bi-level structure, where a Teacher agent suggests tasks to the Student based on the Student’s demonstrated potential. This approach facilitates the learning of simpler tasks initially, enabling the transfer of knowledge to more complex tasks and reducing the number of interactions required to attain a successful policy. However, these methods typically focus on optimizing a curriculum for learning a single policy, which may not be suitable for tasks with extended temporal aspects. In contrast, we propose a Logical Specifications-driven Teacher-Student learning approach that targets the learning of policies for promising automaton transitions, thereby enhancing sample efficiency compared to conventional approaches.

3 Theoretical Framework

Episodic MDP. An episodic labeled Markov Decision Process (MDP) M is a tuple $(\mathcal{S}, \mathcal{A}, P, R, \mathcal{S}_0, \gamma, K, \mathcal{P}, L)$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $P(s'|s, a)$ denotes the transition probability of reaching state $s' \in \mathcal{S}$ from $s \in \mathcal{S}$ using action $a \in \mathcal{A}$, $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, \mathcal{S}_0 is the initial state distribution, $\gamma \in [0, 1]$ is the discount factor, K is the maximum number of interactions in any episode, \mathcal{P} is a set of predicates, and $L : \mathcal{S} \rightarrow 2^{\mathcal{P}}$ is a labeling function that maps a state $s \in \mathcal{S}$ to a subset of predicates that are true in that state. During each interaction, the agent observes the current state s and selects

an action a according to its policy function $\pi(a|s, \theta)$ with parameters θ . The agent transitions to a new state $s' \in \mathcal{S}$ with probability $P(s' | s, a)$. The agent’s goal is to learn an *optimal policy* π^* that maximizes the discounted return $G_0 = \sum_{k=0}^K \gamma^k R(s'_k, a_k, s_k)$ until the end of the episode, which occurs after at-most K interactions.

High level specification language: In our framework, we adopt the specification language SPECTRL to articulate reinforcement learning tasks (Jothimurugan, Alur, and Bastani 2019). A specification ϕ in SPECTRL is a logical formula applied to trajectories, determining whether a given trajectory $\zeta = (s_0, s_1, \dots)$ successfully accomplishes a desired task. Mathematically, ϕ can be depicted as a function $\phi : \mathcal{Z} \rightarrow \mathbb{B}$, where $\mathbb{B} = \{\text{TRUE}, \text{FALSE}\}$ and \mathcal{Z} is the set of all trajectories.

Formally, a specification is defined over a set of *atomic predicates* \mathcal{P}_0 . Each $p \in \mathcal{P}_0$ is associated with a function $f_p : \mathcal{S} \rightarrow \mathbb{B}$. The agent’s MDP state s satisfies p (denoted by $s \models p$) when $f_p(s) = \text{True}$ (in other words, $p \subseteq L(s)$).

The set of *predicates* \mathcal{P} comprises conjunctions and disjunctions of atomic predicates \mathcal{P}_0 . A predicate $b \in \mathcal{P}$ follows the grammar $b ::= p \mid (b_1 \wedge b_2) \mid (b_1 \vee b_2)$, where $p \in \mathcal{P}_0$. Each predicate $b \in \mathcal{P}$ corresponds to a function $f_b : \mathcal{S} \rightarrow \mathbb{B}$ defined naturally over Boolean logic.

The syntax of SPECTRL specifications is given by

$$\phi ::= \text{achieve } b \mid \phi_1 \text{ ensuring } b \mid \phi_1; \phi_2 \mid \phi_1 \text{ or } \phi_2,$$

where $b \in \mathcal{P}$. Here, *achieve* and *ensuring* correspond to ‘eventually’ and ‘always’ operators in temporal logic. Each specification ϕ corresponds to a function $f_\phi : \mathcal{Z} \rightarrow \mathbb{B}$, and $\zeta \in \mathcal{Z}$ satisfies ϕ (denoted $\zeta \models \phi$) if $f_\phi(\zeta) = \text{TRUE}$. The SPECTRL semantics for a finite trajectory ζ of length t are:

$$\zeta \models \text{achieve } b \quad \text{if } \exists i \leq t, s_i \models b \text{ (or } b \subseteq L(s_i)) \quad (1)$$

$$\zeta \models \phi \text{ ensuring } b \quad \text{if } \exists i \leq t, s_i \models b \quad (2)$$

$$\zeta \models \phi_1; \phi_2 \quad \text{if } \exists i < t, \zeta_{0:i} \models \phi_1 \text{ and } \zeta_{i+1:t} \models \phi_2 \quad (3)$$

$$\zeta \models \phi_1 \text{ or } \phi_2 \quad \text{if } \zeta \models \phi_1 \text{ or } \zeta \models \phi_2 \quad (4)$$

Intuitively, the condition (1) signifies that the trajectory should *eventually* reach a state where the predicates b hold true. The condition (2) signifies that the trajectory should satisfy specification ϕ while *always* remaining in states where b holds true. The condition (3) signifies that the trajectory should sequentially satisfy ϕ_1 and then ϕ_2 . The condition (4) signifies that the trajectory should satisfy either ϕ_1 or ϕ_2 . A trajectory ζ satisfies ϕ if there is a t such that the prefix $\zeta_{0:t}$ satisfies ϕ .

Furthermore, each SPECTRL specification ϕ is *guaranteed* to have an equivalent directed acyclic graph (DAG), called an abstract graph. An *abstract graph* $\mathcal{G} = (Q, E, q_0, F, \beta, \mathcal{Z}_{safe}, \kappa)$ is a directed acyclic graph (DAG) with nodes Q , (directed) edges $E \subseteq Q \times Q$, initial node $q_0 \in Q$, final nodes $F \subseteq Q$, subgoal region map $\beta : Q \rightarrow 2^{\mathcal{S}}$ such that for each $q \in Q$, $\beta(q)$ is a subgoal region and *safe trajectories* $\mathcal{Z}_{safe} = \bigcup_{e \in E} \mathcal{Z}_{safe}^e$ where $\mathcal{Z}_{safe}^e \subseteq \mathcal{Z}_f$ denotes the safe trajectories for edge $e \in E$. Intuitively, (Q, E)

is a standard DAG, and q_0 and F define a graph reachability problem for (Q, E) . Furthermore, β and \mathcal{Z}_{safe} connect (Q, E) back to the original MDP M ; in particular, for an edge $e = q \rightarrow q'$, \mathcal{Z}_{safe}^e is the set of trajectories in the MDP M that can be used to transition from $\beta(q)$ to $\beta(q')$ ¹. The function κ labels each edge $e = q \rightarrow q'$ with the predicates b_e (labeled edge denoted as $e := q \xrightarrow{b_e} q'$). The agent transitions from q to q' when the states s_i, s_{i+j} in the agent’s trajectory ζ satisfy $s_i \subseteq \beta(q)$ and $b_e \subseteq L(s_{i+j})$ and $j \geq 0$.

Given a SPECTRL specification ϕ , we can construct an abstract graph \mathcal{G}_ϕ such that, for every trajectory $\zeta \in \mathcal{Z}$, we have $\zeta \models \phi$ if and only if $\zeta \models \mathcal{G}_\phi$. Thus, we can solve the reinforcement learning problem for ϕ by solving the reachability problem for \mathcal{G}_ϕ . As described below, we leverage the structure of \mathcal{G}_ϕ in conjunction with reinforcement learning to do so. In summary, SPECTRL specifications provide a powerful and expressive means to define and evaluate reinforcement learning tasks. It allows users to specify complex conditions and requirements for successful task completion, enabling a nuanced approach to learning from specifications.

Problem Formulation. Given an MDP M with unknown transition dynamics and a SPECTRL formula ϕ representing the high-level task objective, let \mathcal{G}_ϕ be the DAG representing the language of ϕ . Let $\text{Paths}(q, X)$ be the set of all paths in the DAG originating in q and terminating at a node in $X \subseteq Q$. The aim of this work is to learn a set of policies π_i^* , $i = 0, \dots, n-1$, with the following three properties: (i) Following π_0^* results in a trajectory in the MDP that induces a transition from q_0 to some state $q_1 \in Q$ in the DAG, following π_1^* results in a trajectory in MDP that induces a transition from q_1 to some state $q_2 \in Q$ in the DAG, and so on. (ii) The resulting path $q_0 q_1 \dots q_n$ in the DAG terminates at a final node, *i.e.*, $q_n \in F$, with probability greater than a given threshold, $\eta \in (0, 1)$. (iii) The total number of environmental interactions spent in exploring and learning sub-task policies are minimized.

4 Methodology

Sub-task definiton: Given the DAG \mathcal{G}_ϕ representing the language of ϕ , we define a set of sub-tasks based on the edges of the DAG. Intuitively, given any MDP state $s \in \mathcal{S}$ and a DAG node $q \in Q$, a sub-task defined by an edge from node q to $p \in Q$ defines a reach-avoid objective for the agent represented by the SPECTRL formula,

$$\text{Task}(q, p) := \text{achieve}(b_{(q,p)}) \text{ ensuring } \left(\bigwedge_{r \in \text{Sc}(q), r \neq p} \neg b_{(q,r)} \right)$$

where $b_{(q,p)}$ is the propositional formula labeling the edge from q to p in the DAG and $\text{Sc}(q)$ is the set of successors of node q in DAG. For example, in Fig. 1b, the propositional formula labeling the edge from q_0 to q_1 is $b_{(q_0, q_1)} = \neg l \wedge k_1$. When $e = (q, p)$, we use $\text{Task}(e)$ instead of $\text{Task}(q, p)$ and b_e instead of $b_{(q,p)}$ for notational convenience.

¹See DURL (Jothimurugan et al. 2021) for more details

Each sub-task $\text{Task}(q, p)$ defines a problem to learn a policy $\pi_{(q,p)}^*$ such that, given any MDP state $s_0 \in \mathcal{S}$, following $\pi_{(q,p)}^*$ results in a trajectory $s_0 s_1 \dots s_n$ in MDP that induces the path $qq \dots qp$ in the DAG. That is, the agent’s high-level DAG state remains at q until it transitions to p . While constructing the set of sub-tasks, we omit transitions that lead to a ‘sink’ state (from which final states are unreachable).

Given the MDP M with unknown transition dynamics and the SPECTRL objective, ϕ , we first translate ϕ to its corresponding directed acyclic graphical (DAG) representation $\mathcal{G}_\phi = (Q, E, q_0, F, \beta, \mathcal{Z}_{safe}, \kappa)$. Next, we define the set of sub-tasks. For this, we consider the edges that lie on some path in the DAG from q_0 to some node in F . This is because any path that does not visit F leads to a sink state from which the objective cannot be satisfied. Such edges are identified using breadth-first-search (Moore 1959).

LSTS Initialization: The algorithm for *LSTS* is described in Algo. 1. We begin by initializing the following (lines 2-4): (1) Set of: Active Tasks AT , Learned Tasks LT , Discarded Tasks DT ; (2) A Dictionary $\Pi : \text{Task}(e) \rightarrow \pi_e$ that maps the sub-task $\text{Task}(e)$ corresponding to edge e of DAG \mathcal{G}_ϕ to a RL policy π_e ; (3) A Dictionary $Q : \text{Task}(e) \rightarrow \mathbb{R}$ representing the Q-Values for the Teacher agent by mapping $\text{Task}(e)$ to a numerical q-value associated with $\text{Task}(e)$.

Initially, we transform \mathcal{G}_ϕ into an Adjacency Matrix denoted as \mathcal{X} (line 6), then identify the tasks corresponding to all outgoing edges $\bar{E}_{q_0} \subseteq E$ originating from the initial node q_0 (line 7). Satisfying the predicates $b_{(q_0, q_1)} \in \kappa(\bar{E}_{q_0})$ signifies a reachability sub-task M' . In this sub-task, every goal state $s \in \mathcal{S}_f^{M'}$ of M' satisfies the condition $b_{(q_0, q_1)} \subseteq L(s)$. The Student agent is rewarded positively for fulfilling $b_{(q_0, q_1)}$ and negatively at all other time steps. The state, action space, and transition dynamics of M' mirror those of M . To accomplish this sub-task, the Student agent must learn a Reinforcement Learning (RL) policy $\pi_{(q_0, q_1)}^*$ ensuring a visit from q_0 to q_1 with a probability exceeding a predetermined threshold (η). Additionally, the policy must prevent unintended transitions in the Directed Acyclic Graph (DAG); for example, while obtaining *Key*₁, it should not inadvertently obtain *Key*₂.

Teacher-Student learning: To begin, the Teacher Q-Values for all the sub-tasks corresponding to edges in AT (i.e., tasks corresponding to \bar{E}_{q_0}) are set to zero (line 8). We formalize the Teacher’s objective of selecting the most promising sub-task within a *Partially Observable Markov Decision Process* framework (Kaelbling, Littman, and Cassandra 1998). Here, the Teacher lacks access to the complete state of the Student agent but only observes its performance on a sub-task (such as success rate or average returns). The Teacher’s action involves choosing a sub-task $\text{Task}(e) \in \text{AT}$ that the Student agent should prioritize for training. Within this POMDP framework, each action taken by the Teacher is associated with a Q-Value. Generally, higher Q-Values indicate tasks where the Student agent demonstrates greater success, prompting the Teacher to prioritize sampling such tasks more frequently to achieve ϕ

(reaching a goal node) with the fewest overall interactions.

(A) The Teacher agent samples a sub-task $\text{Task}(e) \in \text{AT}$ using the ϵ -greedy exploration strategy (line 10), and (B) The Student agent trains on task $\text{Task}(e)$ using its corresponding RL policy $\Pi[e]$ for a few interactions (line 11). In one Teacher timestep, the Student trains on x environmental interactions. Since the aim of the Teacher is to keep switching the Student agent’s learning progress to a sub-task on which it shows the highest promise, $x \ll$ total number of environmental interactions required by the Student agent to learn a successful RL policy for $\text{Task}(e)$. (C) The Teacher observes the Student agent’s average return g_t on these x interactions, and updates its Q-Value for $\text{Task}(e)$ (line 12) using: $Q[e] \leftarrow \alpha(g_t) + (1 - \alpha)Q[e]$, where α is the Teacher learning rate, g_t is the average Student agent return on $\text{Task}(e)$ at the Teacher timestep t . As learning progresses, g_t increases as t increases, and hence we use a constant parameter α to tackle the non-stationarity of a moving return distribution. Several other algorithms could be used for the Teacher strategy (e.g. Thompson Sampling). Steps (A), (B), (C) continue successively until the policy for *any* $\text{Task}(e) \in \text{AT}$ task converges.

Sub-task convergence criteria: The Student agent’s RL policy for $\text{Task}(q, p)$ is considered to be converged (line 13) if a trajectory ζ produced by the policy produces the transition in the DAG with probability $\Pr_{\zeta \in \mathcal{Z}}[\zeta \text{ satisfies } \text{Task}(q, p)] \geq \eta$ and $\Delta(g_t, g_{t-1}) < \tau$ where η is the expected performance and τ is a small numerical value. In simple terms, a converged policy achieves an average success rate of at least η and shows no further improvement by maintaining consistent average returns. Similar to other RM methods, we assume access to the labeling function L to check if a trajectory ζ satisfies the formula $b_{(q,p)}$ by verifying if the final state s of the trajectory meets the condition $b_{(q,p)} \subseteq L(s)$. It’s worth noting that the sub-goal regions may overlap, meaning the same state s could satisfy propositions for multiple nodes in the Directed Acyclic Graph (DAG). Once a policy for the $\text{Task}(q, p)$ converges, we add $\text{Task}(q, p)$ to the set of Learned Tasks LT and remove it from the set of Active Tasks AT (line 14).

Discarding unpromising sub-tasks: On learning a successful policy for the $\text{Task}(q, p)$ (transition $q \xrightarrow{b_{(q,p)}} p$), we find the sub-tasks to be discarded (line 15). We find the sub-tasks corresponding to edges that: (1) lie before p in a path from q_0 to any $q \in F$, and, (2) do not lie in a path to $q \in F$ that does not contain p . Intuitively, if we already have a set of policies that can generate a successful trajectory to reach the node p , we do not need to learn policies for sub-tasks that ultimately lead *only* to p (e.g., in Fig. 1 if we have successful policies for $\text{Task}(q_0, q_2)$ and $\text{Task}(q_2, q_3)$, we can discard $\text{Task}(q_0, q_1)$ and $\text{Task}(q_1, q_3)$). We add all such sub-tasks to the set of Discarded Tasks DT (line 16). As an extension, in the limit, an optimal policy can be found by not completely discarding such sub-tasks, but rather biasing away from them so that they would still be explored.

Algorithm 1: *LSTS* ($\mathcal{G}_\phi, M, \eta, x$)

Output: Set of learned policies : Π^* ,

Edge-Policy Dictionary: \mathcal{P}

```

1: Placeholder Initialization:
2: Sets of: Active Tasks (AT)  $\leftarrow \emptyset$ ;
   Learned Tasks (LT)  $\leftarrow \emptyset$ ;
   Discarded Tasks (DT)  $\leftarrow \emptyset$ 
3: Edge-Policy Dictionary  $\Pi : \text{Task}(e) \rightarrow \pi$ 
4: Teacher Q-Value Dictionary:  $Q : \text{Task}(e) \rightarrow -\infty$ 
5: Algorithm:
6:  $\mathcal{X} \leftarrow \text{Adjacency\_Matrix}(\mathcal{G}_\phi)$ 
7:  $\text{AT} \leftarrow \text{AT} \cup \{\text{Tasks}(\mathcal{X}[q_0])\}$ 
8:  $\forall \text{Task}(e) \in \text{AT} : Q[e] = 0$ 
9: while True do
10:  $e \leftarrow \text{Sample}(Q)$ 
11:  $\Pi[e], g \leftarrow \text{Learn}(M, \mathcal{G}_\phi, e, x, \mathcal{P})$ 
12:  $\text{Update\_Teacher}(Q, e, g)$ 
13: if Convergence( $Q, e, g, \eta$ ) then
14:    $\Pi^* \leftarrow \Pi^* \cup \Pi[e]$ ;
    $\text{LT} \leftarrow \text{LT} \cup \{\text{Task}(e)\}$ ;
    $\text{AT} \leftarrow \text{AT} \setminus \{\text{Task}(e)\}$ 
15:    $\text{Tasks}(\overline{E}_{DT}) \leftarrow \text{Discarded\_Tasks}(\mathcal{X}, e)$ 
16:    $\text{DT} \leftarrow \text{DT} \cup \text{Tasks}(\overline{E}_{DT})$ 
17:    $\text{Tasks}(\overline{E}_{AT}) \leftarrow \text{Next\_Tasks}(\mathcal{X}, e, \text{DT})$ 
18:   if  $|\text{Tasks}(\overline{E}_{AT})| = 0$  then
19:     break
20:   end if
21:    $\forall \text{Task}(\bar{e}) \in \text{Tasks}(\overline{E}_{AT}) : Q[\bar{e}] = 0$ 
22:    $\text{AT} \leftarrow \text{AT} \cup \text{Tasks}(\overline{E}_{AT})$ 
23: end if
24: end while
25: return  $\Pi^*, \Pi$ 

```

Traversing in the DAG until ϕ satisfied: Next, we identify the subsequent set of tasks $\text{Tasks}(\overline{E}_{AT})$ within the DAG to include in the AT set (line 17). This is determined by pinpointing sub-tasks corresponding to all outgoing edges from p . Given that the edge $e_{q,p}$ represents the transition $q \xrightarrow{b_{(q,p)}} p$, we have learned policies that can produce a trajectory leading to a state where $b_{(q,p)}$ holds true. $\text{Tasks}(\overline{E}_{AT})$ consists of $\mathcal{X}[p] \setminus \text{DT}$ (where ‘ \setminus ’ denotes set-minus), which indicates sub-tasks corresponding to outgoing edges from p not included in the DT set.

Once $\text{Tasks}(\overline{E}_{AT})$ is identified, we assign Teacher Q-values of 0 to all $\text{Task}(\bar{e}) \in \text{Tasks}(\overline{E}_{AT})$ to ensure the Teacher samples these tasks (line 21). In our episodic scenario, each episode initiates from a state $s \sim S_0$ where the propositions for q_0 are true. If the current sampled sub-task is $\text{Task}(p, r)$, the agent follows a trajectory utilizing learned policies from Π^* to reach a state where the propositions for reaching DAG node p are true (i.e., $s \in \beta(p)$). Subsequently, the agent endeavors to learn a separate policy for $\text{Task}(p, r)$.

These aforementioned steps (lines 9-24) continue iteratively until $|\text{Tasks}(\overline{E}_{AT})| = \emptyset$. This signifies that there

Approach	# Interactions (Mean \pm SD)	Success Rate (Mean \pm SD)
<i>LSTS</i>	$(2.72 \pm 0.31) \times 10^6$	0.96 ± 0.02
<i>LSTS</i> ^{ct}	$(2.45 \pm 0.25) \times 10^6$	0.95 ± 0.01
DiRL ^c	$(4.06 \pm 0.37) \times 10^6$	0.95 ± 0.03
DIRL	$(5.47 \pm 0.40) \times 10^6$	0.94 ± 0.01
QRM	5×10^7	0.05 ± 0.04
GSRS	5×10^7	0 ± 0
TSCL	5×10^7	0 ± 0
LFS	5×10^7	0 ± 0

Table 1: Table comparing #interactions & success rate. *LSTS* (highlighted) outperformed all baselines

are no further tasks to incorporate into our sampling strategy, and we have arrived at a node $q \in F$. Hence, we exit the while loop (line 19) and return the set of learned policies Π^* along with the edge-policy dictionary Π (line 25). From Π and Π^* , we obtain an ordered list of policies $\Pi_{list}^* = [\pi_{(q_1, q_2)}, \pi_{(q_2, q_3)}, \dots, \pi_{(q_{k-1}, q_k)}]$ such that sequentially following $\pi \in \Pi_{list}^*$ generates trajectories satisfying the SPECTRL objective ϕ

Guarantee: Given the ordered list of policies Π_{list}^* , we can generate a trajectory ζ in the task M with $\Pr_{\zeta \in \mathcal{Z}}[\zeta \text{ satisfies } \phi] \geq \eta$ (Details in Appendix B).

5 Experimental Results

We aim to answer the following questions: (Q1) Does *LSTS* yield sample efficient learning compared to state-of-the-art baselines? (Q2) After reaching a sub-task goal state, can we sample a new sub-task to continue training and improve sample efficiency? (Q3) Does *LSTS* yield sample efficient learning for complex robotic tasks with partially observable or continuous control settings? (Q4) How does *LSTS* scale to complex search-and-rescue scenarios?

5.1 LSTS - Gridworld Domain

To answer (Q1), we evaluated *LSTS* on a Mini-grid (Chevalier-Boisvert, Willems, and Pal 2018) inspired domain with the SPECTRL objective:

$$\phi_f^{grid} := (\text{achieve}(k_1) \text{ or } \text{achieve}(k_2); \text{achieve}(d); \text{achieve}(g)) \text{ensuring}(\neg l) \quad (5)$$

where k_1, k_2, d, g, l represent the atomic propositions *Key*₁, *Key*₂, *Door*, *Goal*, *Lava* respectively. The environment and its ϕ are given in Fig. 1. The agent needs to *pick-up any* of the *Keys* before heading to the *Door*. After *toggling* the *Door* open, the agent needs to reach the *Goal*. At all times, the agent needs to avoid the *Lava*. For our experiments, we assume an episodic setting where an episode ends if the agent touches the *Lava* object, reaches the *Goal* or exhausts the number of allocated interactions.

The problem is complex as the agent needs to perform a series of high-level correct actions to satisfy ϕ_f^{grid} . The action space of the agent consists of three navigational actions:

move forward, *rotate left* and *rotate right*. The agent can also perform: *pick-up* action that adds a *Key* to the agent’s inventory if the agent is facing the *Key*, *drop* places the *Key* in the next grid if grid is empty and *Key* is present in the inventory, and, a *toggle* action for the *Door* (closed \leftrightarrow open) only if the agent is holding the *Key*. In this experiment, we assume a fully-observable setting where the environmental state is an image-like low-level encoding of the state. For each cell in the grid, the encoding returns an integer that describes the item occupying the grid, along with any additional information (e.g., the *Door* can be open or closed).

We employ PPO (Schulman et al. 2017) for the Student RL agent, which works for discrete and continuous action spaces. We consider a standard actor-critic architecture with 2 convolutional layers followed by 2 fully connected layers. For *LSTS*, the reward function is sparse. The agent gets a reward of $(1 - 0.9^{\frac{\text{interactions taken}}{\text{interactions allocated}}})$ (where $\text{interactions allocated} = 100$) if it achieves the goal in the sub-task, and a reward of 0 otherwise.

Baselines: We compare our *LSTS* method with six baseline approaches: learning from scratch (LFS), Reward Machine-based (RM) baselines: GSRs (Camacho et al. 2018), QRM (Icarte et al. 2018); and Compositional RL from Logical Specifications (DiRL) (Jothimurugan et al. 2021). All the baselines are implemented using the RL algorithm (PPO), described above. GSRs assigns reward inversely proportional to the distance from the RM goal node. QRM employs a separate Q-function for each node in the RM, and DiRL uses Dijkstra’s algorithm (edge cost is the average RL policy success rate) to guide the agent in choosing a path from the specification graph. For the fifth baseline, we modify DiRL such that instead of manually specifying a limit on the number of interactions, which needs to be fine-tuned to suit the task, we stop learning a sub-task once it has reached the convergence criteria defined in Section 4. We call this modified baseline as DiRL^c. The sixth baseline (TSCL (Matiisen et al. 2020)) follows a curriculum learning strategy where the Teacher samples most promising task without the use of any automaton to guide the learning progress of the agent. (More details in Appendix C)

The results in Table 1 (averaged over 10 trials) show that *LSTS* reaches a successful policy quicker compared to all baselines. *LSTS*^{ct} is a modified version of *LSTS* and is described below. Our custom baseline, DiRL^c is more sample-efficient than DiRL, and both outperform other baselines, which do not learn a meaningful policy. We performed an unpaired t-test (Kim 2015) to compare *LSTS* against the best performing baselines at the end of 10^7 training steps and we observed statistically significant results (95% confidence). Thus, *LSTS* not only achieves a better success rate, but also converges faster (statistical significance result details in Appendix D). Time-to-threshold metric is defined as the difference in number of interactions between two approaches to reach a desired performance (Narvekar et al. 2020). The time-to-threshold between *LSTS* and the best-performing baseline DiRL^c is 1.34×10^6 interactions for 95% success rate.

LSTS^{ct} (LSTS + Cont. Training) - Gridworld Domain

In *LSTS*, while learning a policy for Task(q, p), we reinitialize the environment to a random initial environmental state $s \sim \mathcal{S}_0$ once the agent reaches a state where the propositions ($b_{(q,p)}$) hold true. To answer the question Q2, instead of resetting the environment after reaching such a state where $b_{(q,p)}$ hold true, we let the Teacher agent sample a task (let’s say Task(p, r)) from the set $\mathcal{X}[p] \setminus \text{DT}$, where \mathcal{X} is the adjacency matrix for the graph, and DT is the set of Discarded Tasks, as defined in Algo. 1. This helps the agent continue its training by attempting to learn a policy $\pi_{(p,r)}$ for Task(p, r) while simultaneously learning a separate policy $\pi_{(q,p)}$ for Task(q, p). If the agent fails to satisfy Task(p, r), we reset the environment to state $s \sim \mathcal{S}_0$. Otherwise, the agent continues its training until its trajectory satisfies the high-level objective ϕ . We call this approach *LSTS*^{ct} (Detailed algo in Appendix A). Results in Table 1 demonstrate that this approach improves sample efficiency by reducing the number of interactions required to learn a successful policy for the gridworld task, with a time-to-threshold metric of 2.7×10^5 interactions as compared to *LSTS*.

5.2 LSTS and LSTS^{ct} - Robotic Domains

To answer (Q3), we test *LSTS* and *LSTS*^{ct} on two simulated robotic environments with high interaction cost. The task in Fig. 2a has the following SPECTRL objective:

$$\phi_f^{\text{navigation}} := (\text{achieve}(\text{Key}_1) \text{ or } \text{achieve}(\text{Key}_2); \text{achieve}(\text{Goal})) \text{ ensuring } (\neg \text{Lava}) \quad (6)$$

In this task, the agent (a simulated TurtleBot) needs to collect any of the keys (yellow blocks) present in a $[3m, 3m]$ continuous environment before reaching the goal position (gray block). At all times, the agent needs to avoid the lava object (red wall) present in the center. The move forward (backward) action causes the robot to move forward (backward) by $0.1m$ and the robot rotates by $\pi/8$ radians with each rotate action. The pick-up and drop actions have effects similar to the gridworld domain. The robotic domain is more complex as objects can be placed at continuous locations. The agent receives an ego-centric image view of the environment (top-right corner of Fig. 2a), which makes the task partially observable in nature and more complex to get a successful policy. The RL agent is described in Sec. 5.1.

The second environment (Fig. 2c) consists of a simulated robotic arm pushing two objects to their target locations (Gallouédec et al. 2021) with the SPECTRL formula:

$$\phi_f^{\text{manipulation}} := \text{achieve}(p_1) ; \text{achieve}(p_2) \quad (7)$$

where p_1 and p_2 are the atomic propositions for ‘push-object-1’, ‘push-object-2’. The robot has continuous action parameters for moving the arm and a binary gripper action (close/open). An episode begins with the two objects randomly initialized on the table, and the robotic arm has to push these two objects to its final location. The agent receives its current end-effector pose, positions and velocities of the two objects, and the desired goal position for the two objects. For this task, we use the Deep Deterministic Policy Gradient with Hindsight Experience Replay

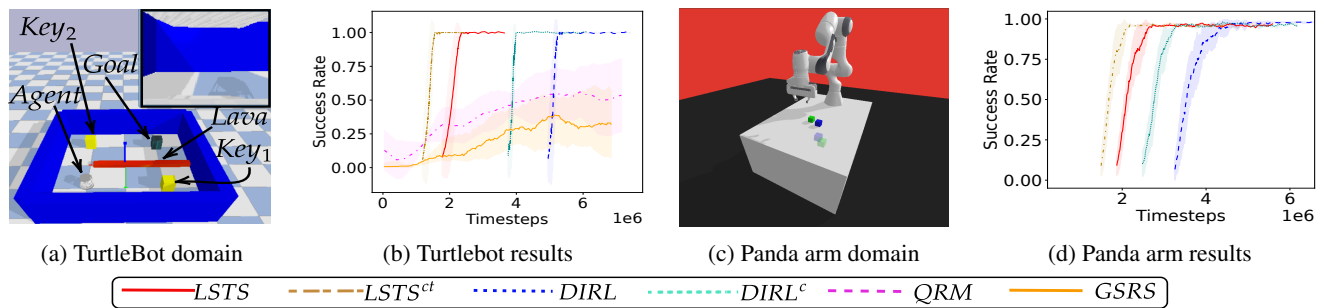


Figure 2: Learning curves (Averaged over 10 trials) for the two robotic domains.

(DDPG-HER) (Andrychowicz et al. 2017) as our RL algorithm. DDPG-HER is implemented using the OpenAI baselines (Dhariwal et al. 2017). Both the robotic domains were modeled using PyBullet (Coumans and Bai 2021), and the reward structure for both the RL agents was sparse, similar to the one described in Sec. 5.1. The learning curves in Fig. 2 have an offset on the x-axis to account for the interactions in the initial sub-tasks before moving on to the final task in the specification, signifying strong transfer. For both domains, *LSTS* outperforms all the baselines in terms of learning speed. *LSTS^{ct}* further speeds-up learning for both the robotic domains. The time-to-threshold between *LSTS* and the best performing baseline (our custom implementation) *DIRL^c*, is 2×10^6 for the TurtleBot domain and 5×10^5 for the Panda arm domain.

5.3 LSTS - Search and Rescue task

To demonstrate how *LSTS* performs when the plan length becomes deeper, we evaluated *LSTS* on a complex urban Search and Rescue domain with multi-goal objectives. In this domain, the agent acts in a grid setting where it needs to perform a series of sequential sub-tasks to accomplish the final goal of the task. The agent needs to open a door using a key, then collect a fire extinguisher to extinguish the fire, and then find and rescue stranded survivors. The order in which these individual sub-goals such as opening the door, rescuing the survivors, and extinguishing the fire are achieved does not matter. A fully-connected graph \mathcal{G}_ϕ generated using the above mentioned high-level states consists of 24 distinct DAG paths. This is a multi-goal task as the agent needs to find the key to open the door, then extinguish fire and rescue survivors to reach the goal state (details in Appendix F). The results in Table 2 (averaged over 10 trials) show that *LSTS* reaches a successful policy quicker compared to the *LFS*, *GSRS*, *QRM* and *TSCL*. The overall number of interactions to learn a set of successful policies for satisfying the high-level goal objective are higher compared to the door key experiment because of the additional complexity of task. We observe that *LSTS* is able to accommodate the task and learn RL policies that satisfy the high-level goal objective.

6 Conclusion

We proposed *LSTS*, a framework for dynamic task sampling for RL agents using the high-level SPECTRL objective

Approach	# Interactions (Mean \pm SD)	Success Rate (Mean \pm SD)
<i>LSTS</i>	$(8.61 \pm 0.12) \times 10^6$	0.87 ± 0.04
<i>LFS</i>	5×10^7	0 ± 0
<i>GSRS</i>	5×10^7	0.05 ± 0.04
<i>QRM</i>	5×10^7	0.05 ± 0.04
<i>TSCL</i>	5×10^7	0 ± 0

Table 2: Table comparing #interactions & success rate for the Search and Rescue domain.

coupled with the Teacher-Student learning strategy. Our approach formalizes each transition of the DAG as a sub-task and identifies promising transitions in the DAG that have the potential to satisfy the high-level objective sample efficiently. Through experiments, we demonstrated that *LSTS* accelerates learning when the high-level objective is correct and available, converging to a desired success rate quicker as compared to other curriculum learning and automaton-guided RL baselines. Furthermore, *LSTS* does not rely on a human-guided dense reward function.

Limitations & Future Work: Our future plans involve incorporating scenarios where obtaining a precise SPECTRL specification is challenging. If the provided specification does not contain any feasible paths, then the algorithm will not be able to invent new feasible paths. A possible solution to this limitation would involve modifying the reward function for each sub-task so that the agent learns policies that enable the agent to reach new high-level states that are not represented in the high-level specification. As an extension, we would like to explore biasing away from sub-tasks rather than completely discarding them once the target node is reached, so in the limit, optimal policies can be obtained.

Acknowledgements

A portion of this work was conducted in the Multimodal Learning, Interaction, and Perception (MuLIP) Lab at Tufts University, University of Texas at Austin, Georgia Tech Research Institute, and the University of Colorado Boulder, with support from the Air Force Research Lab under contract FA875022C0501.

References

- Afzal, M.; Gambhir, S.; Gupta, A.; Trivedi, A.; and Velasquez, A. 2023. LTL-Based Non-Markovian Inverse Reinforcement Learning. In *Proc. of the 2023 Intl. Conf. on Autonomous Agents and Multiagent Systems*, 2857–2859.
- Alur, R.; Bansal, S.; Bastani, O.; and Jothimurugan, K. 2022. A framework for transforming specifications in reinforcement learning. In *Principles of Systems Design*. Springer.
- Andrychowicz, M.; Wolski, F.; Ray, A.; Schneider, J.; Fong, R.; Welinder, P.; McGrew, B.; Tobin, J.; and Pieter Abbeel, W., Zaremba. 2017. Hindsight experience replay. *NeurIPS*.
- Cai, M.; Aasi, E.; Belta, C.; and Vasile, C.-I. 2023. Overcoming Exploration: Deep Reinforcement Learning for Continuous Control in Cluttered Environments From Temporal Logic Specifications. *IEEE RAL*, 8(4): 2158–2165.
- Camacho, A.; Chen, O.; Sanner, S.; and McIlraith, S. A. 2018. Non-Markovian rewards expressed in LTL: Guiding search via reward shaping. In *GoalsRL, a workshop collocated with ICML/IJCAI/AAMAS*.
- Chevalier-Boisvert, M.; Willems, L.; and Pal, S. 2018. Minimalistic Gridworld Environment for Gymnasium.
- Coumans, E.; and Bai, Y. 2021. PyBullet, a Python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>. Accessed: 2023-04-02.
- De Giacomo, G.; and Vardi, M. Y. 2013. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI'13*, 854–860. Association for Computing Machinery.
- Dhariwal, P.; Hesse, C.; Klimov, O.; Nichol, A.; Plappert, M.; Radford, A.; Schulman, J.; Sidor, S.; Wu, Y.; and Zhokhov, P. 2017. OpenAI Baselines.
- Diuk, C.; Cohen, A.; and Littman, M. L. 2008. An object-oriented representation for efficient reinforcement learning. In *25th ICML*, 240–247.
- Gallouédéc, Q.; Cazin, N.; Dellandréa, E.; and Chen, L. 2021. panda-gym: Open-Source Goal-Conditioned Environments for Robotic Learning. *4th Robot Learning Workshop: Self-Supervised and Lifelong Learning at NeurIPS*.
- Grzes, M. 2017. Reward shaping in episodic RL.
- Hammond, L.; Abate, A.; Gutierrez, J.; and Wooldridge, M. 2021. Multi-agent reinforcement learning with temporal logic specifications. *arXiv preprint arXiv:2102.00582*.
- Icarte, R. T.; Klassen, T.; Valenzano, R.; and McIlraith, S. 2018. Using reward machines for high-level task specification and decomposition in reinforcement learning. In *ICML*.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2022. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research*, 73: 173–208.
- Jothimurugan, K.; Alur, R.; and Bastani, O. 2019. A Composable Specification Language for Reinforcement Learning Tasks. In *Advances in NeurIPS*, volume 32, 13041–13051.
- Jothimurugan, K.; Bansal, S.; Bastani, O.; and Alur, R. 2021. Compositional reinforcement learning from logical specifications. *NeurIPS*.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1): 99–134.
- Kim, T. K. 2015. T test as a parametric statistic. *Korean journal of anesthesiology*, 68(6): 540–546.
- Lattimore, T.; Hutter, M.; and Sunehag, P. 2013. The sample-complexity of general reinforcement learning. In *ICML*.
- Matiisen, T.; Oliver, A.; Cohen, T.; and Schulman, J. 2020. Teacher-Student Curriculum Learning. *IEEE Trans. Neural Networks Learn. Syst.*, 31(9): 3732–3740.
- Moore, E. F. 1959. The shortest path through a maze. In *Proc. of the Intl. Symposium on the Theory of Switching*, 285–292. Harvard University Press.
- Narvekar, S.; Peng, B.; Leonetti, M.; Sinapov, J.; Taylor, M. E.; and Stone, P. 2020. Curriculum Learning for Reinforcement Learning Domains: A Framework and Survey. *JMLR*, 21: 1–50.
- Neary, C.; Xu, Z.; Wu, B.; and Topcu, U. 2021. Reward Machines for Cooperative Multi-Agent Reinforcement Learning. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 934–942.
- Oudeyer, P.-Y.; and Kaplan, F. 2009. What is intrinsic motivation? A typology of computational approaches. *Front. in neurorobotics*, 6.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *CoRR*.
- Shukla, Y.; Gao, W.; Sarathy, V.; Velasquez, A.; Wright, R.; and Sinapov, J. 2024. LgTS: Dynamic Task Sampling using LLM-generated sub-goals for Reinforcement Learning Agents. In *23rd Intl. Conf. on Autonomous Agents and Multiagent Systems*.
- Shukla, Y.; Kulkarni, A.; Wright, R.; Velasquez, A.; and Sinapov, J. 2023. Automaton-Guided Curriculum Generation for Reinforcement Learning Agents. In *Proc. of the 33rd Intl. Conf. on Automated Planning and Scheduling*.
- Shukla, Y.; Thierauf, C.; Hosseini, R.; Tatiya, G.; and Sinapov, J. 2022. ACuTE: Automatic Curriculum Transfer from Simple to Complex Environments. In *21st Intl. Conf. on Autonomous Agents and Multiagent Systems*, 1192–1200.
- Szepesvári, C. 2004. Shortest path discovery problems: A framework, algorithms and experimental results. In *AAAI*.
- Toro Icarte, R.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2018. Teaching multiple tasks to an RL agent using LTL. In *Autonomous Agents and MultiAgent Systems*.
- Velasquez, A.; Bissey, B.; Barak, L.; Beckus, A.; Alkhouri, I.; Melcer, D.; and Atia, G. 2021. Dynamic automaton-guided reward shaping for monte carlo tree search. In *Proc. of the AAAI Conf. on Artificial Intelligence*.
- Xu, Z.; and Topcu, U. 2019. Transfer of temporal logic formulas in reinforcement learning. In *IJCAI: Proc. of the Conf.*, volume 28, 4010. NIH Public Access.