Inductive Generalization in Reinforcement Learning from Specifications

Vignesh Subramanian[1][0009−0004−1088−6981], Rohit Kushwah[2], Subhajit Roy[2][0000−0002−3394−023X], Suguman Bansal[1][0000−0002−0405−073X]

[1] School of Computer Science, Georgia Institute of Technology, USA
{vignesh, suguman}@gatech.edu
[2] Department of Computer Science and Engineering, Indian Institute of Technology Kanpur, India
{krohitk, subhajit}@cse.iitk.ac.in

Abstract. We present a novel *inductive generalization framework* for RL from logical specifications. Many interesting tasks in RL environments have a natural inductive structure. These *inductive tasks* have similar overarching goals but they differ inductively in low-level predicates and distributions. We present a generalization procedure that leverages this inductive relationship to learn a higher-order function, a *policy generator*, that generates appropriately *adapted* policies for instances of an inductive task in a zero-shot manner. An evaluation of the proposed approach on a set of challenging control benchmarks demonstrates the promise of our framework in generalizing to unseen policies for long-horizon tasks.

1 Introduction

Formal methods community has contributed strongly to *reinforcement learning (RL)*, especially from formal specifications [1, 7, 9, 10, 13, 14, 20, 25, 38, 40]. These techniques may not provide strong guarantees. In fact, their inability to offer rigorous guarantees has been proven [2, 39]. Nevertheless, these methods provide a principled approach for handling learning over long-horizon tasks.

Generalization remains one of the fundamental challenges in RL. While RL agents can achieve impressive performance on individual tasks, they often struggle to transfer learned behaviors to even slightly modified scenarios. Most RL approaches lack formal mechanisms to capture and exploit structural relationships between tasks, instead relying on implicit generalization through neural network function approximation. This often leads to superficial generalization that fails to capture deeper task similarities. Recent work has attempted to address these challenges through meta-learning and goal-conditioned learning, but developing RL algorithms that can generalize remains an open challenge.

While the challenge of overall generalization is too large to address, this work presents a novel notion of generalization, which we call *inductive generalization*. This is based on leveraging inductive similarities between tasks to generalize. Inductive relationships are fundamental to computational tasks because they capture how complex behaviors can be built from simpler ones through systematic

transformation. They appear naturally whenever tasks exhibit natural recursion or iteration. Our insight is that once we understand how to transform from step i to step $i+1$, we can systematically generalize to handle arbitrarily many steps. This naturally arising pattern is particularly evident in robotics and control tasks, where physical constraints often impose regular structure. By formalizing these inductive relationships, we can move beyond treating each task instance as independent and instead leverage their inherent structural connections to enable systematic generalization.

To this end, we present a *logic-guided approach to inductive generalization* in RL. We use logical specifications to encode a class of *inductive tasks* which comprises of several similar tasks that can be enumerated from each other using an inductive relationship. We require that these tasks have identical logical structure but differ only in the low-level details of predicate values and/or environment parameters. Next, we leverage their similar structures to design a generalizable RL algorithm.

Formally, an inductive task is given by a tuple $\mathsf{R} = (\mathsf{R}_0, \mathsf{update\_pred}, \mathsf{update\_init})$ where R_0 is a *base task* given as a temporal logical specification over given predicates and $\mathsf{update\_pred}$ and $\mathsf{update\_init}$ define inductive updates to the predicates of the specification and environment parameters, respectively, to be applied to the base task repeatedly, so as to generate a family of tasks $\mathsf{R}_1, \mathsf{R}_2, \ldots$ and such. Intuitively, one can think of an inductive task to represent a complex task involving iterations as follows:

```
base_task = \task_0 // Encodes the base spec. and envt. conditions
current_task = base_task
repeat
    next_task = update_pred(current_task), update_init(current_task)
    current_task = next_task
```

For instance, in Figure 1 the robot is required to transport the source pile of boxes to a target pile. This complex task can be decomposed into a series of inductively related tasks: for the i-th instance, pick the topmost box from a height of i in *Source* and place it at the top of the *Target* pile at height $(h-i)$ (where h is the total number of blocks). Here, the 0-th task instance forms the base task, and the update functions change the location of the topmost block in the source and target piles. Observe that all task instances in an inductive task have identical logical structure. They only differ in the instantiations of the low-level predicates and environment variables.

Our goal is to leverage this inductive relationship between tasks to design generalizable RL algorithms to learn policies for each task instance. Concretely, the question we ask is: if trained on a few task instances of an inductive task, can we obtain policies for the remaining tasks in a *zero-shot* manner?

Such generalization is difficult, in general. To solve this problem, we hypothesize that *inductively-related task instances may have inductively-related policies*. Based on this hypothesis, we attempt to learn an *inductive relationship* between the policies of simpler task instances to extract a *policy generator*: a higher-order

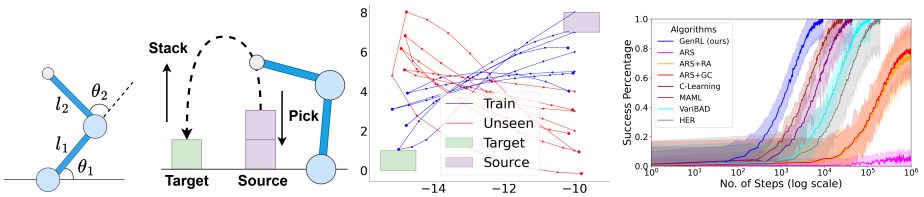

(a) Agent dynamics (b) Inductive Task (c) Trajectories (d) Sample Complexity

Fig. 1: Tower Destacking: The task is to pick boxes from *Source* and stack it on *Target*.

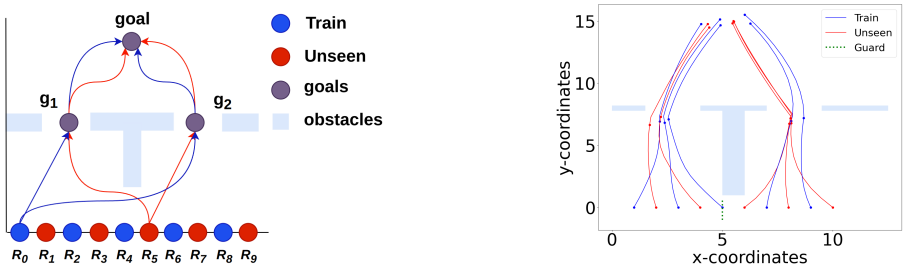

(a) Inductive Task (b) Trajectories

Fig. 2: Choice: visit either g_1 or g_2, then visit goal; task instances differ in initial state distribution.

function that returns an *adapted* policy for a given inductive task instance. Figure 1(c) shows the trajectories of the pickup head with $h = 8$ blocks: we trained a policy generator for the robot on picking and placing the first four blocks (shown in blue); The robot could complete the whole task, with adapted policies from the learned policy generator for the unseen task instances, i.e. pick-n-place of the bottom four blocks are shown in red. We see that the policy generator lends *significant adaptability* to the robot to control its θ_1 and θ_2, as the trajectories of the task instances are quite different.

However, our hypothesis may not always hold. It is possible that despite the task being inductive, the policies are not immediately inductive. The motivating example from Figure 2 illustrates this complication. Figure 2 illustrates an *inductive task* in a 2D Cartesian plane: in a task instance, the agent is initially located in one of the blue or red regions marked R_k. The goal is to *visit the region* marked goal, *after visiting* one of the intermediate regions g_1 or g_2, while *always avoiding* the obstacles shown in light blue. The task is inductive on the initial position: the $(k+1)$-th task can be defined in terms of the k-th task, by shifting the initial location to the right by c units.

However, the policies are not inductive: there is a task R_k such that its policy needs to route through g_1 but the policy of task R_{k+1} must route through g_2

(eg. R_4 and R_5). Yet, we may be able to *classify* the task instances into multiple groups, such that all tasks in each group is *inductive* (eg. $\{\mathsf{R}_0, \ldots, \mathsf{R}_4\}$ and $\{\mathsf{R}_5, \ldots, \mathsf{R}_9\}$). Our policy generator learns such *branches* such that the task instances on the same decision of the branch have inductive policies.

The benchmark environments we utilize in this work are particularly well-suited for evaluating inductive generalization capabilities. Our tasks span a range from simple reachability in 2D environments to complex robotic manipulation scenarios, all unified by their inherent inductive structure. These environments feature continuous state and action spaces, long-horizon planning requirements, and varying degrees of physical constraints, which are challenging even for RL without generalization.

We summarize our contributions: (a). We introduce a framework to learn inductively generalizable policies for long-horizon tasks. This comprises formalizing the notion of inductively-related tasks based on their logical specification and describing the generalization problem as learning a higher-order policy generator (Section 3). (b). We describe a procedure to learn a neural policy generator by leveraging the inductive relationship between task instances (Section 4-Section 5). (c). We perform an empirical evaluation of our inductive framework for generalization in learning unseen tasks in complex, long-horizon specifications in continuous environments, popular control environments, and robotic pick-n-place tasks. Our evaluation demonstrates the promise of our inductive approach (Section 6) as we are able to show that our approach outperforms mature policy-gradient state-of-the-art generalizable RL algorithms in their ability to generalize to unseen tasks and sample complexity.

2 Preliminaries

2.1 Markov Decision Process (MDP)

The environment in RL is given by a Markov Decision Process (MDP) $\mathcal{M} = (S, A, P, \eta)$ with continuous states $S \subseteq \mathbb{R}^n$, continuous actions $A \subseteq \mathbb{R}^m$, transitions $P(s, a, s') = p(s' \mid s, a) \in \mathbb{R}_{\geq 0}$ (i.e., probability density of transitioning from state s to state s' upon taking action a), and initial states $\eta : S \to \mathbb{R}_{\geq 0}$ (i.e., $\eta(s)$ is the probability of the initial state being s).

Let $\mathcal{Z}$ denote the set of all trajectories. A *trajectory* $\zeta \in \mathcal{Z}$ is either an infinite sequence $\zeta = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \cdots$ or a finite sequence $\zeta = s_0 \xrightarrow{a_0} \cdots \xrightarrow{a_{t-1}} s_t$ where $s_i \in S$ and $a_i \in A$. A subtrajectory of ζ is a subsequence $\zeta_{\ell:k} = s_\ell \xrightarrow{a_\ell} \cdots \xrightarrow{a_{k-1}} s_k$. We let $\mathcal{Z}_f$ denote the set of finite trajectories. A (deterministic) *policy* $\pi : \mathcal{Z}_f \to A$ maps a finite trajectory to a fixed action.

Crucially, in RL we assume that the transition probabilities of the MDP are unknown. Hence, the MDP is accessed by sampling only. Concretely, given a policy π, we can sample a trajectory by sampling an initial state $s_0 \sim \eta(\cdot)$, and then iteratively taking the action $a_i = \pi(\zeta_{0:i})$ and sampling a next state $s_{i+1} \sim p(\cdot \mid s_i, a_i)$.

2.2 SPECTRL Specification Language and their Abstract Graphs

We express RL tasks using the logical specification language SPECTRL [18]. Every SPECTRL specification can be expressed as an *abstract graph* which can be used to design scalable compositional algorithms for RL from logical specifications [19].

A SPECTRL specification is defined over a set of *atomic predicates* $\mathcal{P}_0$ that ground environment states, where every $p \in \mathcal{P}_0$ is associated with a function $[\![p]\!] : S \to \mathbb{B} = \{\texttt{true}, \texttt{false}\}$; we say a state s *satisfies* p (denoted $s \models p$) if and only if $[\![p]\!](s) = \texttt{true}$. For $b \in \mathcal{P}$, the syntax of SPECTRL is: $\phi ::= \texttt{achieve}\ b \mid \phi_1\ \texttt{ensuring}\ b \mid \phi_1; \phi_2 \mid \phi_1\ \texttt{or}\ \phi_2$. Each specification ϕ corresponds to a function $[\![\phi]\!] : \mathcal{Z} \to \mathbb{B}$, and we say $\zeta \in \mathcal{Z}$ satisfies ϕ (denoted $\zeta \models \phi$) if and only if $[\![\phi]\!](\zeta) = \texttt{true}$. Intuitively, '`achieve` ' and ' `ensuring` ' are reachability and safety goals, respectively. ';' and '`or`' refer to sequencing and disjunction, respectively. Letting ζ be a finite trajectory of length t, this function is defined by

$$
\begin{aligned}
&\zeta \models \texttt{achieve}\ b && \text{if } \exists\ i \le t,\ s_i \models b \\
&\zeta \models \phi\ \texttt{ensuring}\ b && \text{if } \zeta \models \phi \text{ and } \forall\ i \le t,\ s_i \models b \\
&\zeta \models \phi_1; \phi_2 && \text{if } \exists\ i < t,\ \zeta_{0:i} \models \phi_1 \text{ and } \zeta_{i+1:t} \models \phi_2 \\
&\zeta \models \phi_1\ \texttt{or}\ \phi_2 && \text{if } \zeta \models \phi_1 \text{ or } \zeta \models \phi_2.
\end{aligned}
$$

Abstract Graph. An *abstract graph* of a SPECTRL specification is a DAG-like structure in which vertices represent sets of states (called subgoal regions) and edges represent sets of MDP trajectories that can be used to transition from the source to the target vertex without violating safety constraints.

Definition 1. An *abstract graph* $\mathcal{G} = (U, E, u_0, F, \beta, \mathcal{Z}_{\text{safe}})$ is a directed acyclic graph (DAG) with vertices U, (directed) edges $E \subseteq U \times U$, initial vertex $u_0 \in U$, final vertices $F \subseteq U$, subgoal region map $\beta : U \to 2^S$ such that for each $u \in U$, $\beta(u)$ is a subgoal region, and *safe trajectories* $\mathcal{Z}_{\text{safe}} = \bigcup_{e \in E} \mathcal{Z}^e_{\text{safe}} \cup \bigcup_{f \in F} \mathcal{Z}^f_{\text{safe}}$, where $\mathcal{Z}^e_{\text{safe}} \subseteq \mathcal{Z}$ denotes the safe trajectories for edge $e \in E$ and $\mathcal{Z}^f_{\text{safe}} \subseteq \mathcal{Z}$ denotes the safe trajectories for final vertex $f \in F$.

Intuitively, (U, E) is a DAG, and u_0 and F define a graph reachability problem for (U, E). Furthermore, β and $\mathcal{Z}_{\text{safe}}$ connect (U, E) back to the original MDP $\mathcal{M}$; in particular, for an edge $e = u \to u'$, $\mathcal{Z}^e_{\text{safe}}$ is the set of safe trajectories in $\mathcal{M}$ that can be used to transition from $\beta(u)$ to $\beta(u')$.

A trajectory $\zeta = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \cdots \xrightarrow{a_{t-1}} s_t$ in $\mathcal{M}$ satisfies the abstract graph $\mathcal{G}$ (denoted $\zeta \models \mathcal{G}$) if there is a sequence of indices $0 = k_0 \le k_1 < \cdots < k_\ell \le t$ and a path $\rho = u_0 \to u_1 \to \cdots \to u_\ell$ in $\mathcal{G}$ such that (a). $u_\ell \in F$, (b). for all $z \in \{0, \ldots, \ell\}$, we have $s_{k_z} \in \beta(u_z)$, (c). for all $z < \ell$, letting $e_z = u_z \to u_{z+1}$, we have $\zeta_{k_z:k_{z+1}} \in \mathcal{Z}^{e_z}_{\text{safe}}$, and (d). $\zeta_{k_\ell:t} \in \mathcal{Z}^{u_\ell}_{\text{safe}}$. The first two conditions state that the trajectory should visit a sequence of subgoal regions corresponding to a path from the initial vertex to some final vertex, and the last two conditions state that the trajectory is composed of subtrajectories that are safe according to $\mathcal{Z}_{\text{safe}}$.

The *edge policy* π_e for an edge $e = u \to u'$ is one that safely transitions from a state in $\beta(u)$ to a state in $\beta(u')$. Given edge policies Π along with a path

$\rho = u_0 \to u_1 \to \cdots \to u_k = u$ in $\mathcal{G}$, the *path policy* π_ρ navigates from $\beta(u_0)$ to $\beta(u)$. In particular, π_ρ executes $\pi_{u_j \to u_{j+1}}$ (starting from $j = 0$) until reaching $\beta(u_{j+1})$, after which it increments $j \leftarrow j + 1$ (unless $j = k$). Learning an optimal policy for SPECTRL is reduced to learning an optimal path policy from the initial to final vertex. This gives rise to a natural compositional learning approach that first learns edge policies and then returns the path policy with the maximum probability of reaching a final vertex [19].

3 Inductive Tasks

We begin by describing *inductive tasks*. These appear naturally in several scenarios, as shown in Figures 1-2.

Notation. An *RL task* is given by the tuple (ϕ, η) where ϕ is a SPECTRL specification and η is the initial state distribution in the MDP. We say a trajectory $\zeta = s_0 \dots s_t$ satisfies an RL task (ϕ, η), denoted $\zeta \models (\phi, \eta)$, if $s_0 \sim \eta$ and $\zeta \models \phi$, I.e., ζ begins in a state sampled from η and ζ satisfies ϕ.

An *inductive task* is a family of *RL tasks* that demonstrate the same overarching structure but differ inductively in the low-level details. I.e., an *inductive task* is given by a set of enumerable RL tasks such that the $(i+1)$-th task builds on the i-th task by updating the predicates in the specification and/or the MDP initial distribution. Formally,

Definition 2. *Let $\mathcal{P}$ and $D(S)$ denote the sets of predicates and state distributions in an MDP, respectively. Let $\phi(P)$ denote a SPECTRL specification defined over predicates $P \subseteq \mathcal{P}$. Then, an* inductive task *is given by* $\mathsf{R} = (\mathsf{R}_0, \mathsf{update\_pred}, \mathsf{update\_init})$ *where RL task* $\mathsf{R}_0 = (\phi(\mathcal{P}_0), \eta_0)$ *is the* base task*,* $\mathsf{update\_pred} : \mathcal{P} \mapsto \mathcal{P}$ *is the* predicate update function*, and* $\mathsf{update\_init} : \mathcal{D}(S) \mapsto \mathcal{D}(S)$ *is the* initial distribution update function*. The enumerable* task instances *in* R *are given by* $\mathsf{R}_0 = (\phi(\mathcal{P}_0), \eta_0)$ *and* $\mathsf{R}_{i+1} = (\phi(\mathcal{P}_{i+1}), \eta_{i+1})$ *for* $i > 0$ *where* $\mathcal{P}_{i+1} = \{\mathsf{update\_pred}(p) \mid p \in \mathcal{P}_i\}$ *and* $\eta_{i+1}(s) = \eta_i(\mathsf{update\_init}(s))$.

We denote the i-th task instance R_i by (ϕ_i, η_i) and refer to task instances R_i and R_{i+1} as *adjacent*.

Motivating Example #1. For Figure 1, the inductive task is formalized as: For $j \in \{0, \dots, h\}$, let the predicates $\mathsf{source\_j}$ and $\mathsf{target\_j}$ denote the location of the block at height j in the source and target tower, respectively; let $\eta\_\mathsf{source\_j}$ be a distribution around the block at height j in the source tower. The base task R_0 is given by

$$((\texttt{achieve}\ (\mathsf{target\_0})); (\texttt{achieve}\ (\mathsf{source\_(h-1)}))), \eta\_\mathsf{source\_h}).$$

The predicate update function updates predicates $\mathsf{source\_j}$ and $\mathsf{target\_j}$ to $\mathsf{source\_(j-1)}$ and $\mathsf{target\_(j+1)}$, resp. The initial distribution update function updates $\eta\_\mathsf{source\_j}$ to $\eta\_\mathsf{source\_(j-1)}$. Then, the j-th task instance

$$\mathsf{R}_j = ((\texttt{achieve}\ (\mathsf{target\_j})); (\texttt{achieve}\ (\mathsf{source\_(h-j-1)}))), \eta\_\mathsf{source\_h-j}).$$

Motivating Example #2. Choice tasks from Figure 2 is an inductive task that updates the initial state distribution, by shifting to the right by constant units for adjacent task instances. They are stack of l levels, where each task R_k requires reaching a goal $goal_i$ while avoiding the obstacle obs, either through the (sub)goal g_{i1} or g_{i2},

$$\begin{aligned}&(\texttt{achieve}\ (\texttt{reach}\ (g_{i1})\ \texttt{or}\ \texttt{reach}\ (g_{i2}));\\ &\quad \texttt{achieve}\ (\texttt{reach}\ (goal_i)))^{l}\\ &\texttt{ensuring}\ (\texttt{avoid}\ (obs))\end{aligned}$$

where, $1 \leq i \leq l$. We use the superscript l to indicate that the enclosed specification is repeated l times. Figure 2a illustrates a task with $l = 1$ and Figure 9c illustrates a task with $l = 2$. The update functions for the initial distribution is defined as $\mathsf{update\_init}(\eta(s)) = \eta(s + (c_1, 0))$, where $c_1 = 1$ unit. Intuitively, this corresponds to shifting the support of the distribution η to the left by c_1 units along the x-axis.

Lemma 1. *For an inductive task* R*, let* $\mathcal{G}_i$ *be the abstract graph of the specification of the i-th task instance* R_i*. Then, all the* $\mathcal{G}_i$*s share a common DAG structure with the same initial and final vertices.*

Proof. The proof follows from the construction of abstract graphs in [19].

This lemma asserts that all task instances within an inductive task have identical logical structures. They differ only in the low-level details of the abstract graph.

4 Generalizable RL for Inductive Tasks.

We define the problem of learning generalizable policies for an inductive task by learning a *policy generator*. The *policy generator* for an inductive task R is a function $\mathbb{G} : \mathsf{R} \to \Pi$, where Π is the set of all policies in the MDP. E.g., the policy generator for tower-destacking from Figure 1 maps the j-th task instance to the policy that displaces the source's $(\mathsf{h} - \mathsf{j})$-th block to the target's j-th block, then returns to the source's $(\mathsf{h} - \mathsf{j} - 1)$-th block by manipulating the motor controls θ_1, θ_2. Note these policies are different for each task instance R_j.

Definition 3 (Learning a Policy Generator). *Given an MDP with unknown transitions, an* inductive task R *and a set of a* training task instances Train *s.t. the base task* $\mathsf{R}_0 \in \mathsf{Train}$*, the problem of* generalizable RL *is to learn a* policy generator $\mathbb{G}^* : \mathsf{R} \to \Pi$ *such that*

$$\mathbb{G}^* \in \arg\max_{\mathbb{G}} \frac{1}{|\mathsf{Train}|} \cdot \sum_{\mathsf{R}_j \in \mathsf{Train}} \Pr_{s_0 \sim \eta_j \zeta \sim \mathcal{D}_{\pi_j, s_0}} [\zeta \models \phi_j, \eta_j] \text{ where the policy } \pi_j = \mathbb{G}^*(\mathsf{R}_j)$$

Then, $\pi_j = \mathbb{G}^*(\mathsf{R}_j)$ *for all* $j \in \mathbb{N}$.

I.e., the policy generator optimizes the policies for all training task instances simultaneously, in an attempt to *generalize*, so as to also derive policies for all task instances not present in Train.

4.1 Learning Policy Generator

We present an overview of our approach to learning a policy generator. Learning a higher-order function such as the policy generator is difficult. To make learning a policy generator feasible, we (a) assume *inductive relations between policies* of task instances that are inductively related, (b) leverage similarity between the structure of inductive tasks (Lemma 1), and (c) leverage compositionality of SPECTRL specifications [19].

We leverage the inductive nature of the inductive tasks to learn the policy generator. We base our work on the following hypothesis: *As two adjacent task instances are related by an inductive relation, there may also exist an inductive relation over the corresponding policies of these tasks.* However, this may not hold for certain tasks (eg. Figure 2). We attempt to overcome this with *compositonality*: instead of learning an inductive policy for the whole task, we divide the task into *subtasks* via the abstract graph, where each edge in the abstract graph corresponds to a subtask.

[19] ensures that a policy for a task instance R_i is given by a path policy in its abstract graph $\mathcal{G}_i$. Lemma 1 informs that the DAG structure of all graphs $\mathcal{G}_i$ are identical, say $\mathcal{G}$. Hence, the policy generator can be viewed as a map from task instances to path policies from initial to final vertex in the same graph $\mathcal{G}$ (with appropriate instantiation for edge policies in each task instance). Hence, we learn an inductive relation between the corresponding edge policies of the abstract graphs. We formulate the problem to learn such inductive relations in Subsection: *Learning an Inductive Relation on Edges*.

Last but not least, the edge policies obtained from the inductive relation will result in multiple path policies for each task instance. We are interested in the policy generator to choose the optimal path policy for each task instance. We ensure this by incorporating *guards* along vertices in the common DAG $\mathcal{G}$ that route each task instance along the optimal path in the DAG (Subsection: *Learning the Policy Generator*).

Learning an Inductive Relation on Edges. This section defines an inductive relation between corresponding edges of the abstract graphs of an inductive task and formulates our approach to learn neural inductive relations.

Let $e = u \to v$ be an edge in the common DAG $\mathcal{G}$ of an inductive task. Let π_i denote the edge policy for the i-th task on the edge e in $\mathcal{G}_i$.

Then, an *inductive relation* between these policies is a function $\Omega : \Pi \to \Pi$ s.t. $\pi_{i+1} = \Omega(\pi_i)$. Thus, given the edge policy π_0 in the base task, the inductive relation Ω can be inductively "unrolled" to construct the edge policy for any instance R_i of an inductive task R. That is,

$$\pi_i = \Omega(\pi_{i-1}) = \Omega(\Omega(\pi_{i-2})) = \cdots = \Omega^i(\pi_0)$$

where Ω^i composes Ω with itself i times.

As learning the inductive relation Ω is difficult, we resort to *polynomial approximation*: we approximate the inductive relation Ω over the policies as an

m-degree polynomial. Polynomial approximations are interesting as any function can be approximated as a polynomial up to an arbitrary precision using the Taylor expansion. This reduces learning Ω on edges to inferring the κ-coefficients $(\kappa_m, \cdots, \kappa_0)$ of an m-degree κ-polynomial. Details below:

Neural Policies. If the policy for the i-th task instance $\pi_i \in \Pi$ is implemented by a neural network with parameter vector $[\pi_i]$, then the m-degree polynomial inductive relation is given by

$$[\pi_{i+1}] = \kappa_m \odot [\pi_i]^m + \kappa_{m-1} \odot [\pi_i]^{m-1} + \cdots + \kappa_0 \tag{1}$$

where, the polynomial coefficients, κ_i, are vectors with the same dimension as $[\pi_i]$; the $\odot$ operator is the Hadamard product (element-wise multiplication) of the coefficient vectors κ_i with the parameter vector (weights and biases) of the policy network π_i, and '+' is element-wise vector addition. Then as described earlier, with π_0 as the base policy with parameters $[\pi_0]$, the inductive relation Ω can be inductively "unrolled" to construct the policy network for π_i. Hence, in this case, we attempt to learn an inductive relation between the parameter vectors of the policy (neural) network of task instances.

Learning the Policy Generator. Next, we describe a policy generator on the common DAG $\mathcal{G}$ between all task instances in R. Given the inductive relation and base policy for every edge in $\mathcal{G}$, our goal is to describe a mapping for task instances in R to a path policy in $\mathcal{G}$, as per the edge policies inferred by Equation 1.

Every path from the initial to the final vertex corresponds to a path policy for the i-th task. Elaborating further, for degree m, let $\kappa^e = (\kappa^e_m, \cdots, \kappa^e_0)$ denote the κ-coefficients on the edge $e \in \mathcal{G}$. Let $\rho = e_1 \cdots e_k$ be a path from the initial to a final vertex. Then, a policy for a task R_i is given by the path policy $\pi_i^{e_1} \cdots \pi_i^{e_k}$ where $[\pi_i^{e_j}] = \Omega^i[\pi_0^{e_j}]$. This requires *selection* of a path policy for each R_i.

We assign *guards* at vertices with multiple outgoing edges in $\mathcal{G}$ such that each vertex routes task instances to a unique outgoing edge, ensuring a unique path for every task instance from initial to a final vertex. Formally, a guard in a vertex maps task instances to the outgoing edges from the vertex.

Then, the policy generator for an abstract graph executes as follows: Given a task instance R_i, it uses the guard conditions to determine its unique path from the initial to a final vertex. It returns the path policy along this path as described above. For example, Figure 2a has two possible paths: via g_1 or g_2. We learn a guard, $(i \leq 4)$, that resolves this branching decision at the init node: a task like R_2 would select pass via g_1 while R_6 will via g_2.

Hence, learning a policy generator for a DAG entails learning the $(m+1)$ κ-coefficients of the m-degree κ-polynomial and a base policy for every edge, along with guard conditions for all vertices with multiple outgoing edges.

Algorithm 1 GenRL$(\mathsf{R}, m, \mathsf{Train})$

$\mathcal{G} \leftarrow \mathsf{CommonDAG}(\mathsf{R})$
while vertex $u \in \mathcal{G}$ is chosen in topological order **do**
  Compute $P(u,i), \mathsf{bestIn}(u,i)$ **for** all $i \in \mathsf{Train}$
  $\eta_i^u \leftarrow \mathsf{InduceDistrbution}(u, \mathsf{bestIn}(u,i))$ **for** all $i \in \mathsf{Train}$
  for edge $e = (u,v) \in \mathsf{OutEdges}(u)$ **do**
    $\pi_0^e \leftarrow \mathsf{LearnBasePolicy}(e, \eta_0^u)$
    $\kappa^e \leftarrow \mathsf{LearnKappaCoefficients}(e, m+1, \pi_0^e, \eta^u, \mathsf{Train})$
$\mathsf{Guard} \leftarrow \mathsf{LearnGuardConditions}(\mathcal{G}, \mathsf{bestIn})$
return κ^e, π_0^e for all edge and Guard for all vertices

5 Algorithm

Algorithm 1 (GenRL) takes as input an inductive task R, the degree m of the κ-polynomial, and a finite set of training task instances Train (we assume $0 \in \mathsf{Train}$) and outputs a *policy generator* for R. As described above, this entails learning a base policy and the $(m+1)$ κ-coefficients for edges and guard conditions for vertices in the common DAG $\mathcal{G}$ (with initial vertex u_0) of the inductive task.

GenRL operates in two phases: (1) learn κ-coefficients and base policy for all edges, and (2) learn guard conditions at vertices with multiple outgoing edges.

In the first phase, GenRL traverses the vertices in $\mathcal{G}$ in topological order. While processing a vertex u, we record the *success probability* of the best probability path from the initial vertex u_0 to u for the i-th task instance in $P(u,i)$. We also record $\mathsf{bestIn}(u,i)$ to be the set of incoming vertices to u that are along a best probability path from u_0 to u for the task instance $i \in \mathsf{Train}$. Then,

$$P(u,i) = \begin{cases} 1 & \text{if } u = u_0 \\ \max\{P(w,i) \cdot p_i^{w \to u} \mid w \to u \in \mathsf{InEdges}(u)\} & \text{if } u \neq u_0 \end{cases}$$

$$\mathsf{bestIn}(u,i) = \begin{cases} \emptyset & \text{if } u = u_0 \\ \arg\max_w \{P(w,i) \cdot p_i^{w \to u} \mid w \to u \in \mathsf{InEdges}(u)\} & \text{if } u \neq u_0 \end{cases}$$

where $p_i^{w \to u}$ is the estimated success probability of edge policy of i-th task instance on edge $w \to u$.

Next, we induce a state distribution η_i^u on vertex u for all task instances $i \in \mathsf{Train}$. $\eta_i^{u_0}$ is given by the initial distribution of the task instance R_i. For all other vertices $u \neq u_0$, the state distributions are induced along the *best probability path* from u_0 to u. To this end, η_i^u is induced from states in $\mathsf{bestIn}(u,i)$ using the leaned edge policies along these incoming edges.

Finally, for all outgoing edges $e = u \to v$, we learn the base policy and $(m+1)$ κ-coefficients. The base policy π_0^e is learned as a neural-network policy using standard RL such that π_0^e maximizes the satisfaction of the edge e for the 0-th task instance. I.e., the rewards are designed to encourage π_0^e to safely transition from an MDP state in u to an MDP state in v for the 0-th task instance.

Algorithm 2 $\mathsf{LearnKappaCoefficients}(e, m, \pi_e^0, \Gamma_u, \mathsf{Train})$
Kappa Training using a modified Augmented Random Search

Initialize $\kappa(m)$
  where $m \leftarrow$ number of kappa in the polynomial template
while κ not converged **do**
  $\delta_{samples} \leftarrow \emptyset$
  for $s = 0$ to $n\_samples$ **do**
    $r_{plus} \leftarrow \emptyset$, $r_{minus} \leftarrow \emptyset$
    $\delta \leftarrow \mathsf{SampleDelta}(\kappa)$
    $\kappa_{plus} \leftarrow \mathsf{PerturbKappa}(\kappa, \delta, \delta_{scale})$
    $\kappa_{minus} \leftarrow \mathsf{PerturbKappa}(\kappa, \delta, -\delta_{scale})$
    for $k = 0$ to $|\mathsf{Train}|$ and task R_i where $i \in \mathsf{Train}$ **do**
      $\pi_{plus} \leftarrow \mathsf{KappaPolicy}(\kappa_{plus}, \pi_e^0, k, m)$
      $r_{plus}[k] \leftarrow \mathsf{Reward}(\pi_{plus}, \mathsf{R}_i)$
      $\pi_{minus} \leftarrow \mathsf{KappaPolicy}(\kappa_{minus}, \pi_e^0, k, m)$
      $r_{minus}[k] \leftarrow \mathsf{Reward}(\pi_{minus}, \mathsf{R}_i)$
    $R_{plus} \leftarrow \mathsf{Score}(r_{plus})$
    $R_{minus} \leftarrow \mathsf{Score}(r_{minus})$
    $\delta_{samples}[s] \leftarrow (\delta, R_{plus}, R_{minus})$
  $\delta_\kappa \leftarrow \mathsf{DeltaUpdate}(\delta_{samples})$
  Update $\kappa \leftarrow \mathsf{PerturbKappa}(\kappa, \delta_\kappa, 1)$
return κ

The κ-coefficients are learned using an adaptation of the ARS (Augmented Random Search) (see Algorithm 2). The κ-coefficients capture an inductive relation between the parameters of the policy networks of adjacent task instances; $[\pi_i^e]$ is the parameter vector for the policy network corresponding to the i-th task instance. Let π_i^e be obtained by unrolling the κ-polynomial on the base policy parameters for all $i \in \mathsf{Train}$. Taking rewards of π_i^e to be based on satisfaction of the edge for the i-th task instance (as done for the base policy), κ-polynomials are trained to optimize the $\mathsf{softmin}$ of the rewards over all training task instances.

The second phase learns the guard conditions (see Algorithm 3): In addition to ensuring the uniqueness of the path, we require that the guards choose an edge such that the resulting path policy has a high probability of satisfaction. To this end, for each edge e, we create a set S^e of task instances such that the e appears on a best probability path to a final vertex for those task instances, using backward DAG traversal and bestIn. Finally, the guard on a vertex u is learned as a multi-task classifier with (a) outgoing edges as the class labels, (b) task instance indices as features, and (c) the dataset consists of data points (i, e) s.t. $i \in S^e$ for all outgoing edges e from u.

Algorithm 3 LearnGuardConditions($\mathcal{G}$, bestIn)
Training a Decision Tree Classifier at every edge where decision making is involved

\\ Creating Decision Sets
Initialize $O(v) \leftarrow$ OutVertices(v) for all $v \in \mathcal{G}$
Initialize $Q \leftarrow F$ \\ Q is a queue
while Q is not empty **do**
  vertex $u \leftarrow Q$.dequeue
  for vertex $v \in$ InVertices(u) **do**
    $O(v)$.remove(u)
    for $i \in$ Train **do**
      if $v \in$ (bestIn(u, i)) **then**
        $\mathcal{D}_e(e)$.append(i) where edge $e = (u \rightarrow v)$
    if $O(v)$ is empty **then**
      Q.enqueue(v)
\\ Learning Guard
dataset $D = \{(x, y) \mid x \in X, y \in Y\}$ where $X \leftarrow$ EnvInputValues(R), $Y \leftarrow$ edge e
Guard $\leftarrow$ TrainDecisionTree(D) where Guard $= (f : X \rightarrow Y)$
return Guard

6 Empirical Evaluation

We evaluate GenRL[1] across diverse environments, demonstrating superior generalization and sample efficiency compared to state-of-the-art approaches. Through experiments spanning navigation, long-horizon tasks, complex decision-making, and control benchmarks, we show that explicitly modeling inductive relationships between tasks provides a fundamentally more effective approach to generalization.

6.1 Experimental Setup

Evaluation Methodology. To rigorously evaluate generalization capabilities, each experiment involves training on a fixed set of task instances (Train) and testing on unseen task instances (Unseen) from the same inductive task family. We estimate the probability of success for each task based on 1000 rollouts. A policy π_i is considered successful on task instance $\mathsf{R}_i : (\phi_i, \eta_i)$ if the rollouts ζ satisfy the specification with probability above δ, specifically $\Pr_{\zeta \sim \mathcal{D}_{\pi_i}}[\zeta \models \phi_i, \eta_i] > \delta$, where we set $\delta = 0.9$. All reported results represent the median of five independent runs with different random seeds. For fair comparison, all methods are evaluated on identical Train and Unseen sets.

In all experiments, we train using a 1-degree κ-polynomial approximation of the inductive relationship. All experiments run on a cluster with Intel Xeon Gold 6226 CPUs (2.7 GHz, 24 cores per node) and 192GB RAM per node.

[1] Complete codebase of GenRL and the experimental setup is available at https://anonymous.4open.science/r/GenRL_Zenith-7EEB/

6.2 Baselines and Comparisons

Baselines are state-of-the-art generalizable RL approaches and GenRL ablations:

- Generalizable RL algorithms across three categories:
  - *Inductive Generalization algorithms.* PSMP [16] leverages inductive structures between tasks but operates as a planning approach with known transition probabilities;
  - *Meta-Learning algorithms.* MAML-Reinforce [11] and VariBAD-A2C [42], which enable rapid adaptation through task-specific representations;
  - *Goal-Conditioned algorithms.* C-Learning [29] and HER-DDPG [4] enable generalization by incorporating goal-states directly into the policy input.
- Ablations of GenRL to analyze component contributions:
  - *Augmented Random Search (ARS) [28]:* This ablation uses the standard ARS algorithm, learning a single policy for all tasks in Train. Unlike GenRL, which trains an inductive policy generator that adapts to different tasks, ARS relies on a single reward signal rather than aggregating multiple task-specific rewards.
  - *ARS + Reward Aggregation (ARS+RA):* This ablation incorporates our reward aggregation mechanism (softmin function across tasks) but still learns a single policy instead of generating task-specific policies.
  - *ARS + Goal Conditioning (ARS+GC):* This ablation extends ARS+RA by providing the task index as additional input to the policy, enabling differentiation between tasks while maintaining a single policy.

6.3 Results and Analysis

Simple Reachability Tasks We begin with simple reachability tasks in a 2D Cartesian plane. The inductive tasks (Figures 3–5) require navigating from initial positions (blue and red dots) to target positions (grey boxes) while avoiding obstacles when present. Task instances are generated by shifting either only the initial position (Figures 4, 5) or both initial and goal positions (Figure 3).

The figures show representative trajectories for training tasks (blue) and unseen tasks (red). Despite the apparent simplicity, most baselines struggle with generalization. While PSMP performs adequately when task variations follow simple shifts (Figure 3), it fails with more complex adaptations (Figures 4, 5). ARS+RA shows stronger performance in controlled variations but lacks consistent generalization across broader conditions.

In contrast, GenRL demonstrates substantially superior generalization by producing custom trajectories adapted to each specific task instance rather than applying nearly identical paths across all variations. For the simplest case (Figure 3), GenRL successfully completes up to 99 unseen tasks after 400 iterations—a remarkable 1650% generalization rate relative to the training set. Even with obstacles (Figure 5), GenRL maintains impressive generalization with 22 successful unseen tasks (367% generalization rate). Successful trajectories (solid lines with •) vastly outnumber failures (dotted lines with ×) for GenRL even on unseen tasks, confirming its robust adaptation capabilities.

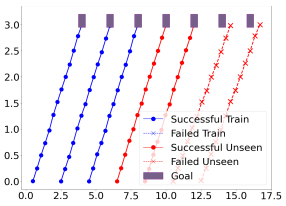

(a) GenRL

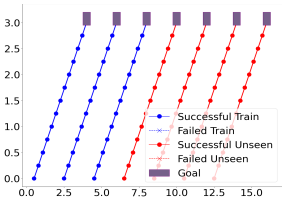

(b) PSMP

Fig. 3: Moving initial and goal distributions

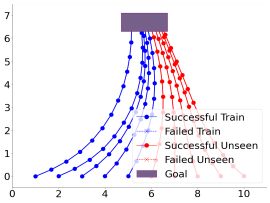

(a) GenRL

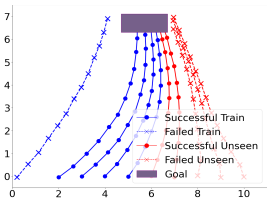

(b) ARS+RA

Fig. 4: Moving initial distribution, stationary goal

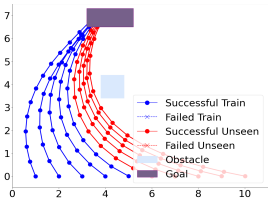

(a) GenRL

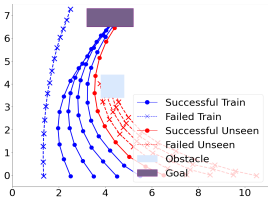

(b) ARS+RA

Fig. 5: Moving initial distribution, stationary goal, with obstacle

Long-Horizon Tasks To evaluate scalability to longer horizons, we designed two classes of long-horizon tasks:

N-Reachability Tasks. These tasks (Figures 6a and 6c) extend simple reachability by requiring the agent to visit N intermediate points sequentially. Both the initial distribution and goal positions shift inductively across task instances.

Figures 6b and 6d show generalization performance as N increases. GenRL maintains consistently high performance across all horizon lengths, successfully generalizing to 5 unseen tasks even in the most complex scenarios ($N = 5$)—an 83% generalization rate. While VariBAD achieves competitive results for $N = 1$, its performance deteriorates dramatically as horizon length increases, highlighting its difficulty with longer-horizon tasks. Other algorithms show poor generalization across the board, with performance declining rapidly as N increases.

Tower-Destacking. We evaluated several variations of a tower-destacking task using a robotic arm in the Reacher environment (Figures 1 and 7). The source tower contains eight blocks, with algorithms trained on the top four blocks and tested on the remaining four.

Figure 8a shows that GenRL achieves generalization performance comparable to sophisticated methods like VariBAD and HER. However, the learning curves in Figure 7 reveal a critical advantage: GenRL reaches optimal performance with an order of magnitude fewer samples than these alternatives. While GenRL converges after approximately 10^4 environment steps, VariBAD requires 10^5 steps, and HER-DDPG needs over 10^6 steps to achieve similar performance. This remark-

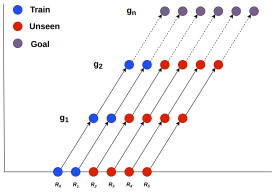

(a) N-Reachability without obstacles (Illustration)

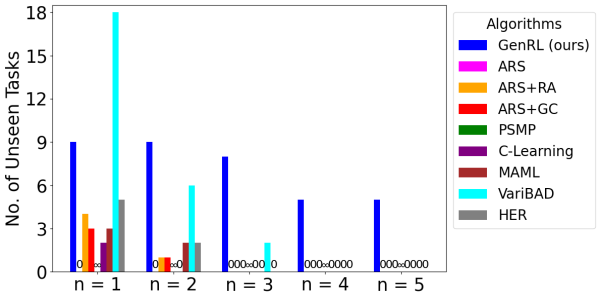

(b) N-Reachability without obstacles (results)

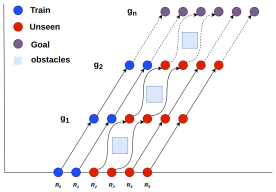

(c) N-Reachability with obstacles (Illustration)

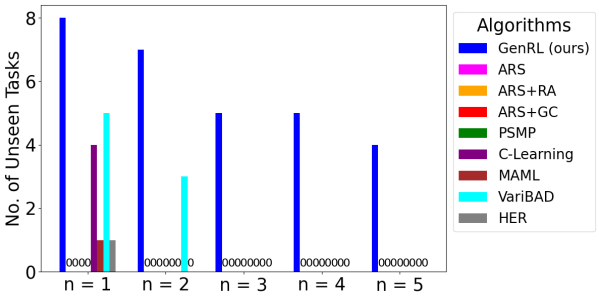

(d) N-Reachability with obstacles (results)

Fig. 6: N-Reachability Tasks: (a) and (c) illustrate the task environments without and with obstacles, respectively. (b) and (d) compare the number of successful unseen tasks for these environments.

able sample efficiency demonstrates the power of explicitly modeling inductive relationships rather than relying on implicit learning through neural networks.

Complex Decision-Making Tasks Some of the most challenging scenarios involve optimal decision-making where the agent must choose between alternative paths based on task parameters. We evaluated GenRL on three "choice" tasks of increasing complexity: (a). Moving initial point with fixed goal (Figure 2) (b). Moving initial point and moving goal (Figure 9a) (c). Two-level choice with moving initial point and moving goal (Figure 9c). These tasks require the agent to select between alternative subgoals (g_1 or g_2) based on the specific task instance. GenRL's unique ability to learn guard conditions enables optimal branching decisions. For example, in Figure 2, GenRL learns the guard predicate: $i \leq 4$, directing tasks with index $i \leq 4$ through the first subgoal and tasks with $i > 4$ through the second subgoal.

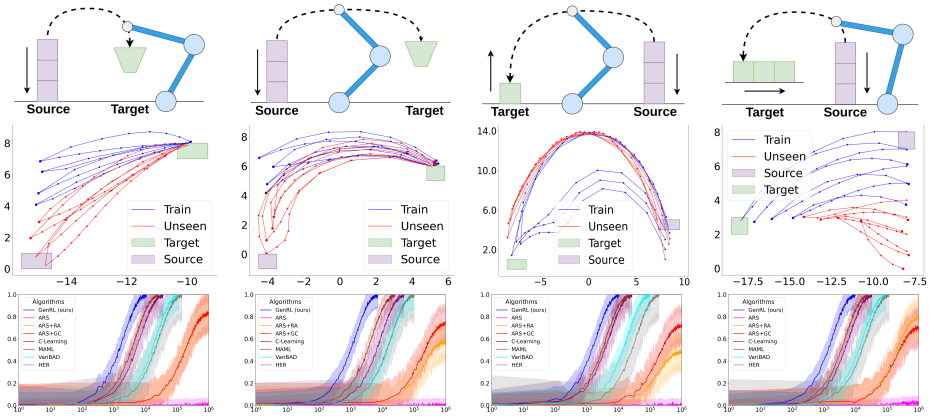

(a) (b) (c) (d)

Fig. 7: Tower-destacking benchmarks on Reacher Environment: Task illustrations, trajectories, and learning curves: (a) Pick and Drop: Same side, (b) Pick and Drop: Opposite side, (c) Pick and Vertical Stack: Opposite side, (d) Pick and Horizontal Stack: Same side. Learning Curve: x-axis denotes the number of samples (steps) and y-axis denotes the average of the estimated probability of success of all tasks in Train. Results are averaged over 5 runs with the cloud indicating the minimum and maximum.

Table 1 shows that GenRL successfully generalizes to 5-7 unseen tasks across these choice benchmarks, achieving up to 117% generalization on the most basic choice task (Figure 2) and 84% generalization on more complex choice variants. Remarkably, no other algorithm demonstrates any meaningful generalization on these tasks, highlighting GenRL's distinctive capability to model complex decision boundaries.

Classical Control Benchmarks To verify GenRL's applicability beyond navigation tasks, we evaluated it on classic OpenAI Gym control benchmarks: Pendulum, Acrobot, and Cartpole. These environments feature inductive variations in physical parameters (mass, length) rather than task specifications.

Figure 8b shows that GenRL achieves superior generalization across all three environments, with particularly impressive results in Pendulum (20 successful unseen tasks, 333% generalization) and Cartpole (14 successful unseen tasks, 233% generalization). Even in the challenging Acrobot environment, GenRL successfully generalizes to 10 unseen tasks (167% generalization) with sufficient training iterations, while competing methods plateau at much lower levels of generalization. This demonstrates that GenRL can effectively model inductive relationships in physical parameters and transfer knowledge between related control tasks.

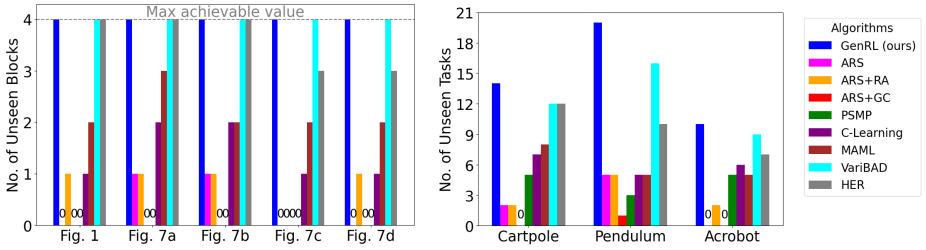

(a) Reacher Benchmarks (b) Classical control benchmarks

Fig. 8: Comparison of the number of successful Unseen task instances. (a) Performance on Reacher benchmarks. (b) Performance on classical Control benchmarks.

Benchmark	Successful Train	Successful Unseen	Learned Guard Predicate
Figure 2	**All**	7	$(i \leq 4)$
Figure 9a	**All**	5	$(i \leq 4)$
Figure 9c	**All**	5	$(i \leq 4)$, $(i \leq 4)$

Table 1: Choice benchmarks: No. of successful unseen tasks for the Car2D Choice benchmarks. **Successful Train** indicates if all training tasks were completed successfully ($|\mathsf{Train}| = 6$), while **Successful Unseen** reports the number of successful instances on unseen tasks. **Learned Guard Predicate** represents the decision index where the agent must choose the optimal goal. For example, when $i \leq 4$, for task R_i where $i \leq 4$, goal g_1 is chosen; otherwise, goal g_2 is chosen. In more complex benchmarks (e.g., Figure 9c with 2 levels), multiple predicates may emerge.

6.4 Key Findings and Analysis

Superior Performance and Stability Across Complexity Scales. GenRL consistently outperforms all baselines across diverse specifications, particularly as task complexity increases. While baseline methods exhibit diminishing returns—evident in N-Reachability tasks where VariBAD's performance drops dramatically from N=1 to N=5, and in choice tasks where baselines fail to generalize at all—GenRL maintains stable performance even for longer horizons and complex branching tasks. This stability, coupled with remarkable sample efficiency (requiring only 10^4 environment steps compared to 10^5-10^6 for policy gradient methods), demonstrates that explicitly modeling inductive relationships is fundamentally more effective than implicit meta-learning or goal-conditioning approaches. GenRL achieves this despite using the simpler Augmented Random Search algorithm without gradients, yet matches or exceeds the performance of sophisticated methods like VariBAD and HER.

This dramatic efficiency advantage demonstrates the power of leveraging logical specifications and inductive structure. For tasks with natural inductive

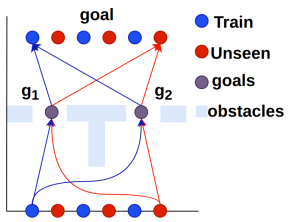

(a) Choice with moving goal - Illustration.

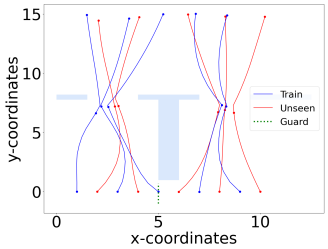

(b) Choice with moving goal - Trajectory.

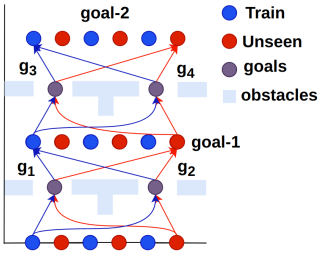

(c) Two levels of choice with moving goal - Illustration.

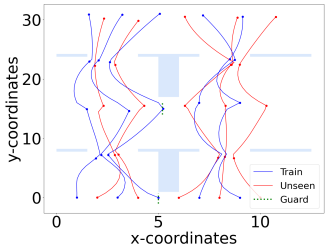

(d) Two levels of choice with moving goal - Trajectory.

Fig. 9: Choice benchmarks task illustration and corresponding trajectories: (a) and (b) illustrate the benchmark, while (c) and (d) show the agent's trajectory in the respective benchmarks.

relationships, explicitly modeling these relationships proves far more effective than attempting to learn them implicitly through meta-learning or policy gradients.

Policy Generator is key to GenRL*'s superior performance.* Our ablation studies with ARS, ARS+RA, and ARS+GC consistently show these variants underperforming compared to the full GenRL framework. The base ARS algorithm has no mechanism to generalize over tasks, resulting in poor performance. Although incorporating Reward Aggregation (ARS+RA) and Goal Conditioning (ARS+GC) provides marginal improvements, these variants still fall far short of GenRL's capabilities.

When ARS leverages GenRL's inductive policy generator framework, it not only overcomes its high sample complexity but significantly enhances performance. This integration enables GenRL to outperform even sophisticated policy gradient methods, demonstrating the substantial advantages of our approach. GenRL's ability to efficiently generate adapted policies for diverse task instances—rather than learning or fine-tuning individual policies—underlies its exceptional sample efficiency and performance.

7 Related Work

Specification-Guided RL. Recent years have seen an emergence of RL from logical specifications [1, 2, 5, 7, 9, 12–15, 17, 19, 20, 25, 26, 33, 38, 40]. Here, the task is expressed using high-level logical specifications rather than as low-level rewards. Logic specifications have received traction due to their ease in expressing complex long-horizon tasks and ability to efficiently scale learning. Prior works have focused primarily on scalability to long-horizon tasks and theoretical guarantees. Ours is the first work to leverage logical specifications specifically for generalization.

Zero-shot generalization. Zero-shot generalization relates to multi-task learning and skill transfer, distilling transferable skills from seen tasks to generalize to unseen ones [22, 30–32]. With logical specifications, existing approaches learn policies for sub-specifications and generalize to their combinations [23, 24, 27, 35, 37]. Our problem is orthogonal: in prior work, the predicate set remains constant while only specifications change. In our setting, both predicates and environment distributions vary between training and unseen tasks. [21] considers changing distributions but with fixed predicates. Reward-based generalizable RL has been explored in [43],[34]. In meta-learning, algorithms like MAML[11] and VariBAD[42] enable rapid adaptation to new tasks by adjusting task-specific parameters and representations. Goal-conditioned algorithms like C-Learning[29] and HER[4] contribute by conditioning policies on desired outcomes, enhancing adaptation to new goals without additional training.

Inductive/Programmatic Approaches to Generalization. Closest to our work is PSMP [16], which learns inductive policies in the planning setting (known MDP) rather than RL (unknown MDP). Despite this advantage, PSMP cannot adapt to different task instances, as it learns a single policy for all instances. Our approach learns a higher-order policy generator that produces specialized policies for each task instance. Programmatic/logic-based policy representations generally demonstrate better generalizability than neural network policies [6, 8, 36, 41]. Non-programmatic policy sketches have also been explored [3]. Our work differs by exploiting the natural inductiveness in task specifications to extract inductive relations for policies and learn a higher-order policy generator.

8 Concluding remarks

While current advances in generalizable RL have focused primarily on making agents more adaptable through sophisticated architectures and training procedures, our work suggests an alternative path forward - one that leverages the inherent structure in specifications to enable systematic generalization. By making this structure explicit through formal specifications and inductive relationships, we not only achieve better generalization but also achieve better sample complexity by an order of magnitude. In doing so, we outperform several mature policy-gradient based state-of-the-art tools for generalization.

Currently, GenRL performs effectively in environments with lower-dimensional action and state spaces. However, its scalability to more complex environments with higher-dimensional spaces remains a challenge. Future work will focus on enhancing the algorithm's capability to handle these more complex scenarios. Additionally, while defining tasks via logical specifications is generally easier than specifying rewards, it still requires considerable effort to design these specifications accurately. To address these, future research will aim at developing more streamlined and user-friendly methods for task specification to make the specification process as lightweight as possible. This will help broaden the applicability of GenRL or any specification-guided learning to a wider range of tasks and environments. Even though the logic and formulation behind our research are principally motivated by a foundational hypothesis and empirically validated for its performance and results, we still not have investigated the possibilities of providing theoretical guarantees and this is also something our future work will focus on.

Bibliography

[1] Aksaray, D., Jones, A., Kong, Z., Schwager, M., Belta, C.: Q-learning for robust satisfaction of signal temporal logic specifications. In: Conference on Decision and Control (CDC). pp. 6565–6570. IEEE (2016)

[2] Alur, R., Bansal, S., Bastani, O., Jothimurugan, K.: 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, pp. 604–624. Springer (2022)

[3] Andreas, J., Klein, D., Levine, S.: Modular multitask reinforcement learning with policy sketches. In: International conference on machine learning. pp. 166–175. PMLR (2017)

[4] Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., Zaremba, W.: Hindsight experience replay. Advances in Neural Information Processing Systems **30** (2017)

[5] Bansal, S.: Specification-guided reinforcement learning. In: International Static Analysis Symposium. pp. 3–9. Springer (2022)

[6] Bastani, O., Pu, Y., Solar-Lezama, A.: Verifiable reinforcement learning via policy extraction. Advances in neural information processing systems **31** (2018)

[7] Brafman, R., De Giacomo, G., Patrizi, F.: LTLf/LDLf non-markovian rewards. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)

[8] Cao, Y., Li, Z., Yang, T., Zhang, H., Zheng, Y., Li, Y., Hao, J., Liu, Y.: GALOIS: boosting deep reinforcement learning via generalizable logic synthesis. Advances in Neural Information Processing Systems **35**, 19930–19943 (2022)

[9] De Giacomo, G., Iocchi, L., Favorito, M., Patrizi, F.: Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications. In: Proceedings of the International Conference on Automated Planning and Scheduling. vol. 29, pp. 128–136 (2019)
[10] Dohmen, T., Perez, M., Somenzi, F., Trivedi, A.: Regular reinforcement learning. In: Gurfinkel, A., Ganesh, V. (eds.) Computer Aided Verification. pp. 184–208. Springer Nature Switzerland, Cham (2024)
[11] Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning. pp. 1126–1135. PMLR (2017)
[12] Hahn, E.M., Perez, M., Schewe, S., Somenzi, F., Trivedi, A., Wojtczak, D.: Omega-Regular Objectives in Model-Free Reinforcement Learning. In: Tools and Algorithms for the Construction and Analysis of Systems. pp. 395–412 (2019)
[13] Hasanbeig, M., Kantaros, Y., Abate, A., Kroening, D., Pappas, G.J., Lee, I.: Reinforcement learning for temporal logic control synthesis with probabilistic satisfaction guarantees. In: Conference on Decision and Control (CDC). pp. 5338–5343 (2019)
[14] Hasanbeig, M., Abate, A., Kroening, D.: Logically-constrained reinforcement learning. arXiv preprint arXiv:1801.08099 (2018)
[15] Icarte, R.T., Klassen, T., Valenzano, R., McIlraith, S.: Using reward machines for high-level task specification and decomposition in reinforcement learning. In: International Conference on Machine Learning. pp. 2107–2116. PMLR (2018)
[16] Inala, J.P., Bastani, O., Tavares, Z., Solar-Lezama, A.: Synthesizing Programmatic Policies that Inductively Generalize. In: Internation Conference on Learning Representations (2020)
[17] Jiang, Y., Bharadwaj, S., Wu, B., Shah, R., Topcu, U., Stone, P.: Temporal-Logic-Based Reward Shaping for Continuing Reinforcement Learning Tasks. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 7995–8003 (2021)
[18] Jothimurugan, K., Alur, R., Bastani, O.: A composable specification language for reinforcement learning tasks. Advances in Neural Information Processing Systems **32** (2019)
[19] Jothimurugan, K., Bansal, S., Bastani, O., Alur, R.: Compositional reinforcement learning from logical specifications. Advances in Neural Information Processing Systems **34**, 10026–10039 (2021)
[20] Jothimurugan, K., Bansal, S., Bastani, O., Alur, R.: Specification-guided learning of Nash equilibria with high social welfare. In: International Conference on Computer Aided Verification. pp. 343–363. Springer (2022)
[21] Jothimurugan, K., Hsu, S., Bastani, O., Alur, R.: Robust Subtask Learning for Compositional Generalization. In: International Conference on Machine Learning (2023)
[22] Kirk, R., Zhang, A., Grefenstette, E., Rocktäschel, T.: A survey of zero-shot generalisation in deep reinforcement learning. Journal of Artificial Intelligence Research **76**, 201–264 (2023)

[23] Kuo, Y.L., Katz, B., Barbu, A.: Encoding formulas as deep networks: Reinforcement learning for zero-shot execution of LTL formulas. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5604–5610. IEEE (2020)
[24] León, B.G., Shanahan, M., Belardinelli, F.: Systematic generalisation through task temporal logic and deep reinforcement learning. arXiv preprint arXiv:2006.08767 (2020)
[25] Li, X., Vasile, C.I., Belta, C.: Reinforcement learning with temporal logic rewards. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3834–3839. IEEE (2017)
[26] Littman, M.L., Topcu, U., Fu, J., Isbell, C., Wen, M., MacGlashan, J.: Environment-Independent Task Specifications via GLTL. arXiv preprint arXiv:1704.04341 (2017)
[27] Liu, J., Shah, A., Rosen, E., Jia, M., Konidaris, G., Tellex, S.: Skill Transfer for Temporal Task Specification. In: CoRL 2023 Workshop on Learning Effective Abstractions for Planning (LEAP) (2023)
[28] Mania, H., Guy, A., Recht, B.: Simple random search of static linear policies is competitive for reinforcement learning. In: Advances in Neural Information Processing Systems. pp. 1805–1814 (2018)
[29] Naderian, P., Loaiza-Ganem, G., Braviner, H.J., Caterini, A.L., Cresswell, J.C., Li, T., Garg, A.: C-learning: Horizon-aware cumulative accessibility estimation. International Conference on Learning Representations (2021)
[30] Oh, J., Singh, S., Lee, H., Kohli, P.: Zero-shot task generalization with multi-task deep reinforcement learning. In: International Conference on Machine Learning. pp. 2661–2670. PMLR (2017)
[31] Sodhani, S., Zhang, A., Pineau, J.: Multi-task reinforcement learning with context-based representations. In: International Conference on Machine Learning. pp. 9767–9779. PMLR (2021)
[32] Sohn, S., Oh, J., Lee, H.: Hierarchical reinforcement learning for zero-shot generalization with subtask dependencies. Advances in neural information processing systems **31** (2018)
[33] Svoboda, J., Bansal, S., Chatterjee, K.: Reinforcement learning from reachability specifications: Pac guarantees with expected conditional distance. In: Forty-first International Conference on Machine Learning (2024)
[34] Taiga, A.A., Agarwal, R., Farebrother, J., Courville, A., Bellemare, M.G.: Investigating Multi-task Pretraining and Generalization in Reinforcement Learning. In: The Eleventh International Conference on Learning Representations (2023)
[35] Vaezipoor, P., Li, A.C., Icarte, R.A.T., Mcilraith, S.A.: LTL2ACTION: Generalizing LTL Instructions for Multi-task RL. In: International Conference on Machine Learning. pp. 10497–10508. PMLR (2021)
[36] Verma, A., Murali, V., Singh, R., Kohli, P., Chaudhuri, S.: Programmatically interpretable reinforcement learning. In: International Conference on Machine Learning. pp. 5045–5054. PMLR (2018)
[37] Xu, D., Fekri, F.: Generalizing LTL Instructions via Future Dependent Options. arXiv preprint arXiv:2212.04576 (2022)