

PG3: Policy-Guided Planning for Generalized Policy Generation

Ryan Yang*, Tom Silver*, Aidan Curtis, Tomás Lozano-Pérez, Leslie Pack Kaelbling

MIT Computer Science and Artificial Intelligence Laboratory
{ryanyang, tslvr, curtisa, tlp, lpk}@mit.edu

Abstract

A longstanding objective in classical planning is to synthesize policies that generalize across multiple problems from the same domain. In this work, we study generalized policy search-based methods with a focus on the score function used to guide the search over policies. We demonstrate limitations of two score functions — policy evaluation and plan comparison — and propose a new approach that overcomes these limitations. The main idea behind our approach, Policy-Guided Planning for Generalized Policy Generalization (PG3), is that a candidate policy should be used to guide planning on training problems as a mechanism for evaluating that candidate. Theoretical results in a simplified setting give conditions under which PG3 is optimal or admissible. We then study a specific instantiation of policy search where planning problems are PDDL-based and policies are lifted decision lists. Empirical results in six domains confirm that PG3 learns generalized policies more efficiently and effectively than several baselines.

1 Introduction

How can we compile a transition model and a set of training tasks into a reactive policy? Can these policies generalize to large tasks that are intractable for modern planners? These questions are of fundamental interest in AI planning (Fikes, Hart, and Nilsson 1972), with progress in *generalized planning* recently accelerating (Srivastava 2011; Bonet and Geffner 2015; Jiménez, Segovia-Aguas, and Jonsson 2019; Rivlin, Hazan, and Karpas 2020).

Generalized policy search (GPS) is a flexible paradigm for generalized planning (Levine and Humphreys 2003; Segovia-Aguas, Jiménez, and Jonsson 2021). In this family of methods, a search is performed through a class of generalized (goal-conditioned) policies, with the search informed by a *score function* that maps candidate policies to scalar values. While much attention has been given to different policy representations, there has been relatively less work on the score function. The score function plays a critical role: if the scores are uninformative or misleading, the search will languish in less promising regions of policy space.

In this work, we propose Policy-Guided Planning for Generalized Policy Generation (PG3), which centers around a new score function for GPS. Given a candidate policy, PG3 solves the set of training tasks *using the candidate to guide planning*. When a plan is found, the agreement between the plan and the policy contributes to the score. Intuitively, if the policy is poor, the planner will ignore its guidance or fail to find a plan, and the score will be poor; if instead the policy is nearly able to solve problems, except for a few “gaps”, the planner will rely heavily on its guidance, for a good score.

PG3 combines ideas from two score functions for GPS: *plan comparison*, which plans on the training tasks and records the agreement between the found plans and the candidate policy; and *policy evaluation*, which executes the policy on the training tasks and records the number of successes. While plan comparison provides a dense scoring over policy space, it works best when the solutions for a set of problems all conform to a simple policy. Absent this property, GPS can overfit to complicated and non-general solutions. Policy evaluation scores are extremely sparse within the policy space, effectively forcing an exhaustive search until reaching a region of the policy space with non-zero scores. However, the policies resulting from this approach are significantly more compact and general. PG3 combines the strengths of plan comparison and policy evaluation to overcome their individual limitations. See Figure 1.

In experiments, we study a simple but expressive class of *lifted decision list* generalized policies for PDDL domains (Mooney and Califf 1995; Khardon 1999; Levine and Humphreys 2003). Since the policies are lifted, that is, invariant over object identities, they can generalize to problems involving new and more objects than seen during learning. We propose several *search operators* for lifted decision list policies, including a “bottom-up” operator that proposes policy changes based on prior planning experience.

This paper makes the following main contributions: (1) we propose PG3 as a new approach to GPS; (2) in a simplified setting, we provide conditions under which PG3 is optimal or admissible for policy search; (3) we propose a specific instantiation of GPS for PDDL domains using lifted decision lists; and (4) we report empirical results across six PDDL domains, demonstrating that PG3 learns policies that generalize to held-out test problems, doing so far more efficiently and effectively than several baselines.

*These authors contributed equally.




Figure 1: (Left) In this work, we study generalized policy search (GPS). Each node represents a candidate policy. (Left middle) A candidate policy for the Forest domain. In this domain, the agent *can* walk along the marked trail (green), but it can also walk in the dirt (brown) as long as it avoids water (blue). The candidate policy, which has the agent walk along the trail, would always succeed, except that the trails sometimes contain rocks (gray), which must be climbed; the policy thus fails when a rock is encountered. This policy is “close” to satisficing (see Definition 1 for a precise meaning), so we would like to give it a good score. Each of the next three columns illustrates a different score function evaluated using two training problems (top and bottom rows). Squares (circles) denote agreement (disagreement) between the candidate policy and plan. (Middle) Policy evaluation gives a trivial score to the candidate policy because it never reaches the goal. (Middle right) Plan comparison gives a poor score because the policy differs substantially from the paths found by a planner, which often leave the marked trail. (Right) We propose PG3, which evaluates the candidate by policy-guided planning, resulting in a good score.

2 Related Work

Policy search has been previously considered as an approach to generalized planning (Levine and Humphreys 2003; Khardon 1999; Jiménez and Jonsson 2015; Segovia-Aguas, Jiménez, and Jonsson 2018). Most prior work either uses *policy evaluation* as a score function, or relies on classical planning heuristics for the search through policy space. One exception is Segovia-Aguas et al. (2021), who extend policy evaluation with goal counting; we include this baseline in our experiments. Because of the limited search guidance and size of the policy space, all of these methods can typically learn only small policies, or require additional problem specifications (Illanes and McIlraith 2019).

Another strategy for generalized planning is to construct a policy, often represented as a finite state machine, from example plans (Levesque 2005; Hu and Levesque 2009; Srivastava et al. 2011a; Winner 2008). Such approaches are successful if the plans all conform to a single compact policy. This assumption is violated in many interesting problems including ones that we investigate here (c.f. *plan comparison*). Other approaches learn *abstractions* from example plans that lend themselves to compact policies (Martin and Geffner 2004; Srivastava et al. 2011b; Bonet and Geffner 2018). While these abstraction-based approaches expand the space of solvable problems, they are still unable to learn policies that deviate dramatically from the example plans.

Recent work has proposed using deep learning to learn a generalized policy (Groshev et al. 2018) or a heuristic function (Rivlin, Hazan, and Karpas 2020; Karia and Srivastava 2020; Gehring et al. 2021). Inspired by these efforts, we in-

clude a graph neural network baseline in experiments. Beyond classical planning in factored and relational domains, learning a goal-conditioned policy is of interest in reinforcement learning and continuous control (Schaul et al. 2015; Pong et al. 2018). Unlike in generalized planning, these settings do not assume a known, structured domain model, and often make very different assumptions about state and action representations. Applying the insights in this paper to these settings is an interesting area for future investigation.

We propose policy-guided planning as a mechanism for *policy learning*. Previous work has also considered the question of how to use a *fixed* (often previously learned) policy to aid planning (Khardon 1999; Yoon, Fern, and Givan 2008).

3 Preliminaries

We begin with a brief review of classical planning and then define the generalized problem setting.

Classical Planning

In AI planning, we are given a *domain* and a *problem*, both often expressed in PDDL (McDermott et al. 1998). We use the STRIPS subset with types and negative preconditions.

A domain is a tuple $\langle P, A \rangle$ where P is a set of *predicates* and A is a set of *actions*. A predicate $p \in P$ consists of a name and an arity. An *atom* is a predicate and a tuple of *terms*, which are either *objects* or *variables*. For example, `on` is a predicate with arity 2; `on(X, Y)` and `on(b1, b2)` are atoms, where X, Y are variables and $b1, b2$ are objects. Terms may be typed. A *literal* is an atom or its negation. An action $a \in A$ consists of a name and

a tuple $\langle \text{PAR}(a), \text{PRE}(a), \text{ADD}(a), \text{DEL}(a) \rangle$, which are the parameters, preconditions, add effects, and delete effects of the action respectively. The parameters are a tuple of variables. The preconditions are a set of literals, and the add and delete effects are sets of atoms over the parameters. A *ground action* is an action with objects substituted for parameters. For example, if `move` is an action with parameters $(?from, ?to)$, then `move(15, 16)` is a ground action. In general, actions may be associated with variable costs; in this work, we focus on *satisficing* planning and assume all costs are unitary.

A domain is associated with a set of problems. A problem is a tuple $\langle O, I, G \rangle$ where O is a set of objects, I is an *initial state*, and G is a *goal*. States and goals are sets of atoms with predicates from P instantiated with objects from O . Given a state S , a ground action $\underline{a} = a(o_1, \dots, o_k)$ with $a \in A$ and $o_i \in O$ is *applicable* if $\text{PRE}(\underline{a}) \subseteq S$.¹ Executing an applicable action \underline{a} in a state S results in a *successor state*, denoted $\underline{a}(S) = (S \setminus \text{DEL}(\underline{a})) \cup \text{ADD}(\underline{a})$. Given a problem and domain, the objective is to find a *plan* $(\underline{a}_1, \dots, \underline{a}_m)$ that *solves* the problem, that is, $\underline{a}_m \circ \dots \circ \underline{a}_1(I) \subseteq G$ and each \underline{a}_i is applicable when executed.

Problem Setting: Generalized Planning

AI planning is traditionally concerned with solving individual planning problems. In *generalized planning*, the objective instead is to find a unified solution to multiple problems from the same domain. Here we are specifically interested in learning a *generalized policy* π , which maps a state S and a goal G to a ground action \underline{a} , denoted $\pi(S, G) = \underline{a}$. This form of π is very general, but our intention is to learn a *reactive* policy, which produces an action with minimal computation, and does not, for example, plan internally. Given a problem $\langle O, I, G \rangle$, the policy is said to *solve* the problem if there exists a plan $(\underline{a}_1, \dots, \underline{a}_m)$ that solves the problem, and such that $\pi(S_i, G) = \underline{a}_i$ for each state S_i in the sequence of successors. In practice, we evaluate a policy for a maximum number of steps to determine if it solves a problem.

Our aim is to learn a policy that generalizes to many problems from the same domain, including problems that were not available during learning. We therefore consider a problem setting with a *training* phase and a *test* phase. During training, we have access to the domain $\langle P, A \rangle$ and a set of *training problems* $\Psi = \{\langle O_1, I_1, G_1 \rangle, \dots, \langle O_n, I_n, G_n \rangle\}$. The output of training is a single policy π . During the test phase, π is evaluated on a set of held-out *test problems*, often containing many more objects than those seen during training. The objective of generalized policy learning is to produce a policy π that solves as many test problems as possible. A policy that solves all problems in the test set is referred to as *satisficing*.

4 Policy-Guided Planning for Policy Search

In this work, we build on *generalized policy search* (GPS), where a search through a set of generalized policies is guided by a *score function*. Specifically, we perform a

¹That is, if all positive atoms in $\text{PRE}(\underline{a})$ are in S , and no negated atoms in $\text{PRE}(\underline{a})$ are in S . We use this shorthand throughout.

greedy best-first search (GBFS), exploring policies in the order determined by SCORE, which takes in a candidate policy π , the domain $\langle P, A \rangle$, and the training problems Ψ , and returns a scalar value, where lower is better (Algorithm 1).

Example 1 (Forest). Forest is a PDDL domain illustrated in Figure 1; see the caption for a description. There is a simple satisficing policy for this domain: walk along the trail, and climb over any rocks that are encountered. This policy is similar in spirit to the classic wall-following policy for mazes. The good-but-imperfect candidate policy illustrated in Figure 1 walks along the trail, but does not climb. When this candidate is considered during GPS, it will spawn many successor policies: one will add `(isRock ?to)` as a precondition to `rule1`; another will create a rule with the `climb` action; among others. Each will then be scored.

The score function has a profound impact on the efficacy of GPS. One possible score function is **policy evaluation**: the candidate policy is executed in each training problem, and the score is inversely proportional to the number of problems solved. A major limitation of this score function is that its outputs are trivial for all policies that do not completely solve any problems, such as the policy above.

Another possible score function is **plan comparison**: a planner is used to generate plans for the training problems, and the candidate policy is scored according to the agreement between the plans and the candidate policy (Algorithm 4). When there are multiple ways to solve a problem, this score function can sharply mislead GPS. For example, plan comparison gives a poor score to the follow-and-climb policy in Example 1, even though the policy is satisficing! This issue is not limited to the Forest domain; similar issues arise whenever goal atoms can be achieved in different orders, or when the same state can be reached through two different paths from the initial state. This phenomenon also arises in “bottom-up” generalized planning (Section 2).

Our main contribution is **Policy-Guided Planning for Generalized Policy Generation (PG3)**. Given a candidate policy, PG3 runs *policy-guided planning* on each training problem (Algorithm 3). Our implementation of policy-guided planning is a small modification to A* search: for each search node that is popped from the queue, in addition to expanding the single-step successors, we roll out the policy for several time steps (maximum 50 in experiments, or until the policy is not applicable), creating nodes for each step. The nodes created by policy execution are given cost 0, encouraging the search to prefer paths generated by the policy. For each training problem, if policy-guided planning returns a plan, we run *single plan comparison* (Algorithm 4) to get a score for the training problem. For each state and action in the plan, the policy is evaluated at that state and compared to the action; if they are not equal, the score is increased by 1. If a plan was not found, the maximum possible plan length ℓ is added to the overall score. Finally, the per-plan scores are accumulated to arrive at a policy score. To accumulate, we use *max* for its theoretical guarantees (Section 4), but *mean* has similar empirical performance.

Intuitively, if a candidate policy is able to solve a problem except for a few “gaps”, policy-guided planning will rely heavily on the policy’s suggestions, and the policy will ob-

GENERALIZED POLICY SEARCH VIA GBFS

```

input: domain  $\langle P, A \rangle$  and training problems  $\Psi$ 
input: search operators  $\Omega$ 
// E.g., an empty decision list
initialize: trivial generalized policy  $\pi_0$ 
// Ordered low-to-high priority
initialize: empty priority queue  $q$ 
// See Algorithm 2
Push  $\pi_0$  onto  $q$  with priority  $\text{SCORE}(\pi_0, P, A, \Psi)$ 
// Repeat until max iters
for  $i = 1, 2, 3, \dots$  do
  Pop  $\pi$  from  $q$  // Best policy in  $q$ 
  for search operator  $\omega \in \Omega$  do
    for  $\pi' \in \omega(\pi, P, A, \Psi)$  do
      Push  $\pi'$  to  $q$  with  $\text{SCORE}(\pi', P, A, \Psi)$ 
  // Policy with the lowest score
return Best seen policy  $\pi^*$ 

```

Algorithm 1: Generalized policy search via GBFS. See Algorithm 2 for the PG3-specific implementation of SCORE.

PG3 SCORE FUNCTION FOR POLICY SEARCH

```

input: candidate policy  $\pi$ 
input: domain  $\langle P, A \rangle$  and training problems  $\Psi$ 
hyperparameter: max plan horizon  $\ell$ 
initialize:  $scores \leftarrow []$ 
for  $\langle O, I, G \rangle \in \Psi$  do
  // See Algorithm 3
  Run POLICYGUIDEDPLAN( $\pi, O, I, G, P, A$ )
  if a plan  $p$  is found then
    // See Algorithm 4
     $score \leftarrow \text{PLANCOMPARISON}(\pi, p, I)$ 
     $scores.append(score)$ 
  else
     $scores.append(\ell)$ 
return  $max(scores)$ 

```

Algorithm 2: Scoring via policy-guided planning (PG3).

tain a high score from PG3. If the candidate policy is poor, policy-guided planning will ignore its suggestions, resulting in a low score. For example, consider again the good-but-imperfect policy from Example 1. Policy-guided planning will take the suggestions of this policy to follow the marked trail until a rock is encountered, at which point the policy becomes stuck. The planner will then rely on the single-step successors to climb over the rock, at which point it will again follow the policy, continuing along the trail until another rock is encountered or the goal is reached.

Limitation. PG3 requires planning during scoring, which can be computationally expensive. Nonetheless, in experiments, we will see that PG3 (implemented in Python) can quickly learn policies that generalize to large test problems.

Theoretical Results

We now turn to a simplified setting for theoretical analysis.

Example 2 (Tabular). Suppose we were to represent policies as tables. Let $\{(S, G) \mapsto \underline{a}\}$ denote a policy that assigns action \underline{a} for state S and goal G . Consider a single search op-

POLICY-GUIDED PLANNING

```

input: policy  $\pi$ 
input: problem  $\langle O, I, G \rangle$  and domain  $\langle P, A \rangle$ 
hyperparameter: max plan horizon  $\ell$ 
// See definition below
return  $A^*$  with GETSUCCESSORS

```

Subroutine GETSUCCESSORS

```

input: state  $S$  // State in search
hyperparameter: max policy execution steps  $k$ 
for ground actions  $\underline{a}$  applicable in  $S$  do
  // Standard successor function
  yield  $\underline{a}(S)$  with cost from domain
// Policy-guided successors
for  $i = 1, 2, \dots, k$  do
  if  $\pi$  is not applicable in  $S$ , stop.
  // Get action from policy
   $\underline{a}_i \leftarrow \pi(S, G)$ 
  // Update current state
   $S \leftarrow \underline{a}_i(S)$ 
  // Yield all encountered states
  yield  $S$  with cost 0

```

Algorithm 3: Policy-guided planning. A helper for the PG3 score function (Algorithm 2). In our experiments with PDDL domains, A^* uses the h_{Add} heuristic (Bonet and Geffner 2001).

SINGLE PLAN COMPARISON

```

input: policy  $\pi$ , plan  $p$ , initial state  $I$ 
initialize:  $score$  to 0 and  $S$  to  $I$ 
for  $a$  in  $p$  do
  if  $\pi(S) \neq \underline{a}$  then
    Add 1 to  $score$ 
     $S \leftarrow \underline{a}(S)$ 
return  $score$ 

```

Algorithm 4: Scoring a policy by comparison to a single plan. A helper for the PG3 score function (Algorithm 2).

erator for GPS, which changes one entry of a policy’s table in every possible way. For example, in a domain with two states S_1, S_2 , one goal G , and two actions $\underline{a}_1, \underline{a}_2$, the GPS successors of the policy $\{(S_1, G) \mapsto \underline{a}_1, (S_2, G) \mapsto \underline{a}_1\}$ would be $\{(S_1, G) \mapsto \underline{a}_2, (S_2, G) \mapsto \underline{a}_1\}$ and $\{(S_1, G) \mapsto \underline{a}_1, (S_2, G) \mapsto \underline{a}_2\}$. Here we are not concerned with generalization; training and test problems are the same, Ψ .

Our main theoretical results concern the influence of PG3 on GPS. Proofs are in Appendix A.

Definition 1 (GPS cost-to-go). *The GPS cost-to-go from a policy π_0 is the minimum k s.t. there exists sequences of policies π_0, \dots, π_k and search operators $\omega_0, \dots, \omega_{k-1}$ s.t. π_k is satisficing, and $\forall j, \pi_{j+1} \in \omega_j(\pi_j)$.*

Note that in the tabular setting, the GPS cost-to-go is equal to the minimum number of entries in the policy table that need to be changed to solve all problems.

Assumption 1. *The heuristic used for A^* search in policy-guided planning (Algorithm 3) is admissible; planning is complete; and all costs in the original problem are unitary.*

Theorem 1. *Under Assumption 1, in the tabular setting (Example 2), if a policy π solves all but one of the problems in Ψ , then the PG3 score for π is equal to the GPS cost-to-go.*

As a corollary, GPS with PG3 will perform optimally under the conditions of Theorem 1, in terms of the nodes expanded by GBFS before reaching a satisficing policy.

Theorem 2. *Under Assumption 1, in the tabular setting (Example 2), PG3 is a lower bound on the GPS cost-to-go.*

These results do not hold for other choices of score functions, including policy evaluation or plan comparison. However, the results do not immediately extend beyond the tabular setting; in general, a single application of a search operator could change a policy’s output on every state, converting a “completely incorrect” policy into a satisficing one. In practice, we expect GPS to lie between these two extremes, with search operators often changing a policy on more than one, but far fewer than all, states in a domain. Toward further understanding the practical case, we next consider a specific instantiation of GPS that will allow us to study PG3 at scale.

Generalized Policies as Lifted Decision Lists

We now describe a hypothesis class of *lifted decision list* generalized policies that are well-suited for PDDL domains (Mooney and Califf 1995; Levine and Humphreys 2003).

Definition 2 (Rule). *A rule ρ for a domain $\langle P, A \rangle$ is a tuple:*

- $\text{PAR}(\rho)$: parameters, a tuple of variables;
- $\text{PRE}(\rho)$: preconditions, a set of literals;
- $\text{GOAL}(\rho)$: goal preconditions, a set of literals;
- $\text{ACT}(\rho)$: the rule’s action, from A ;

with all literals and actions instantiated over the parameters, and with all predicates in P .

As with actions, rules can be *grounded* by substituting objects for parameters, denoted $\underline{\rho} = \rho(o_1, \dots, o_k)$. Intuitively, a rule represents an existentially quantified conditional statement: if there exists some substitution for which the rule’s preconditions hold, then the action should be executed. We formalize these notions with the following two definitions.

Definition 3 (Rule applicability). *Given a state S and goal G over objects \mathcal{O} , a rule ρ is applicable if $\exists(o_1, \dots, o_k)$ s.t. $\text{PRE}(\underline{\rho}) \subseteq S$ and $\text{GOAL}(\underline{\rho}) \subseteq G$, where $\underline{\rho} = \rho(o_1, \dots, o_k)$ and with each $o_i \in \mathcal{O}$.*

Definition 4 (Rule execution). *Given a state S and goal G where rule ρ is applicable, let (o_1, \dots, o_k) be the first² tuple of objects s.t. the conditions of Definition 3 hold with $\underline{\rho} = \rho(o_1, \dots, o_k)$. Then the execution of ρ in (S, G) is $\underline{a} = \text{ACT}(\underline{\rho})$, denoted $\rho(S, G) = \underline{a}$.*

Rules are the building blocks for our main generalized policy representation, the lifted decision list.

Definition 5 (Lifted decision list policy). *A lifted decision list policy π is an (ordered) list of rules $[\rho_1, \rho_2, \dots, \rho_\ell]$. Given a state S and goal G , π is applicable if $\exists i$ s.t. ρ_i is applicable. If π is applicable, the execution of π is the execution of the first applicable rule, denoted $\pi(S, G) = \underline{a}$.*

²First under a fixed, arbitrary ordering of object tuples, e.g., lexicographic, to avoid nondeterminism.

See Appendix E for examples of lifted decision list policies represented using PDDL-like syntax. Compared to PDDL operators, in addition to the lack of effects and the addition of goal preconditions, it is important to emphasize that the rules are *ordered*, with later rules only used when previous ones are not applicable. Also note these policies are *partial*: they are only defined when they are applicable.

Representational Capacity. We selected lifted decision lists because they are able to compactly express a rich set of policies across several domains of interest. For example, repeated application of a lifted decision list policy can lead to looping behavior. Nonetheless, note the absence of numeric values, universal quantifiers, and more sophisticated control flow and memory. PG3 is not specific to lifted decision lists and could be used with richer policy classes.

Generalized Policy Search Operators

Here we present search operators for GPS with lifted decision list policies (Algorithm 1). Recall each operator $\omega \in \Omega$ is a function from a policy π , a domain $\langle P, A \rangle$, and training problems Ψ , to a set of *successor policies* Π' .

Add Rule. This operator adds a new rule to the given policy. One successor policy is proposed for each possible new rule and each possible position in the decision list. Each new rule ρ corresponds to an action with preconditions from the domain, and no goal conditions: for $a \in A$, $\rho = \langle \text{PAR}(a), \text{PRE}(a), \emptyset, a \rangle$. The branching factor for this operator is therefore $|A|(|\pi|+1)$, where $|\pi|$ denotes the number of rules in the given decision list.

Delete Rule. This operator deletes a rule from the given policy. The branching factor for this operator is thus $|\pi|$.

Add Condition. For each rule ρ in the given policy, this operator adds a literal to $\text{PRE}(\rho)$ or $\text{GOAL}(\rho)$. The literal may be positive or negative and the predicate may be any in P . The terms can be any of the variables that are already in ρ . The branching factor for this operator is thus $O(|A||\pi||P|^k)$, where m is the maximum arity of predicates in P and k is the number of variables in ρ .

Delete Condition. For each rule ρ in the given policy, this operator deletes a literal from $\text{PRE}(\rho)$ or $\text{GOAL}(\rho)$. Literals in the action preconditions are never deleted. The branching factor is therefore at most $\sum_{\rho \in \pi} |\text{PRE}(\rho)| + |\text{GOAL}(\rho)|$.

Induce Rule from Plans. This final operator is the most involved; we describe it here at a high level and give details in Appendix B. Unlike the others, this operator uses plans generated on the training problems to propose policy modifications. In particular, the operator identifies a state-action pair that disagrees with the candidate policy, and amends the policy so that it agrees. The mechanism for amending the policy is based on an extension of *triangle tables* (Fikes, Hart, and Nilsson 1972). This operator proposes a single change to each candidate, so the branching factor is 1.

Given these operators, we perform GPS starting with an empty lifted decision list using the operator order: Induce Rule from Plans; Add Condition; Delete Condition; Delete Rule; Add Rule. In Appendix C, we describe two optimizations to improve GPS, irrespective of score function.

Domains	PG3 (Ours)		Policy Eval		Plan Comp		Combo		Goal Count		GNN BC		Random	
	Eval	Time	Eval	Time	Eval	Time	Eval	Time	Eval	Time	Eval	Time	Eval	Time
Delivery	1.00	1.5	0.00	–	0.10	5.8	0.20	2132.4	1.00	27.9	0.40	39.9	0.00	–
Forest	1.00	107.2	1.00	605.1	0.16	51.5	1.00	815.4	0.89	662.4	0.13	–	0.03	–
Gripper	1.00	79.5	0.00	–	0.20	0.8	0.00	–	0.00	–	0.06	–	0.00	–
Miconic	1.00	434.8	0.00	–	0.10	316.6	0.00	–	0.90	2415.3	0.12	–	0.13	–
Ferry	1.00	16.1	0.00	–	0.90	1.8	0.00	–	1.00	5705.2	0.11	11.8	0.00	–
Spanner	1.00	17.7	1.00	6.8	0.56	2.1	1.00	31.2	1.00	5.2	0.37	18.1	0.06	–

Table 1: **Policy learning results.** Eval: fraction of test problems solved by the final learned policy. Time: average wall-clock time (in seconds) required to learn a policy that solves $\geq 90\%$ of test problems, with a missing entry if such a policy was never found. All entries are means across 10 random seeds and 30 test problems per seed, with standard deviations shown in Table 2.

5 Experiments and Results

The following experiments evaluate the extent to which PG3 can learn policies that generalize to held-out test problems that feature many more objects than seen during training.

Experimental Setup. Our main experiments consider the fraction of test problems solved by learned policies. Note that here we are evaluating whether policies are capable of solving the problems on their own; we are not using the policies as planning guidance. Each policy is executed on each test problem until either the goal is reached (success); the policy is not applicable in the current state (failure); or a maximum horizon is exceeded (failure) (Appendix D). All experimental results are over 10 random seeds, where training and test problem instances are randomly generated for each seed. GPS methods are run for 2500 node expansions.

Domains. We use the following domains:

- **Forest:** See Figure 1 for description.
- **Delivery:** An agent must deliver packages to multiple locations while avoiding trap locations. The agent can pick up any number of packages from a home base, move between locations, and deliver a package to a location.
- **Miconic:** In this International Planning Competition (IPC) domain, an agent can move an elevator up or down and board or depart passengers. Each problem requires the agent to transport passengers to their desired floors.
- **Gripper:** In this IPC domain, an agent with grippers can pick, move, and drop balls. Each problem requires the agent to transport balls to their target locations.
- **Ferry:** In this IPC domain, an agent can sail, board, and debark cars. Each problem requires the agent to board cars, sail to their destination, and debark them.
- **Spanner:** In this IPC domain, an agent can move, pick up spanners, and tighten nuts using spanners. Each problem requires the agent to move along the corridor and pick up spanners needed to tighten nuts at the shed at the end of the corridor. The agent cannot move backwards.

See Appendix D for problem counts, sizes, and more details.

Approaches

We evaluate the following methods and baselines:

- **PG3 (Ours):** Our main approach.
- **Policy evaluation:** GPS with policy evaluation.
- **Plan comparison:** GPS with plan comparison. Plans are collected on the training problems using A^* with h_{Add} .

- **Combo:** GPS with policy evaluation, but with ties broken using the plan comparison score function.
- **Goal count:** GPS with a score function that runs the candidate policy on each training problem and counts the number of goal atoms not yet achieved at the end of execution (Segovia-Aguas, Jiménez, and Jonsson 2021).
- **Graph neural network behavior cloning (GNN BC):** Inspired by recent works that use graph neural networks (GNNs) in PDDL domains, this method learns a GNN policy via behavior cloning using plans from the training problems. The architecture and training method are taken from (Silver et al. 2021); see Appendix F.
- **Random:** a trivial baseline that selects among applicable actions uniformly at random on the test problems.

Results and Discussion. See Table 1 for our main empirical results. The results confirm that PG3 is able to efficiently guide policy search toward policies that generalize to larger held-out test problems. Qualitatively, the policies learned by PG3 are compact and intuitive (Appendix E). Baseline performance is mixed, in most cases failing to match that of PG3. The strongest baseline is goal count, which confirms findings from previous work (Segovia-Aguas, Jiménez, and Jonsson 2021); however, even in domains with competitive final evaluation performance, the time required to learn a good policy can substantially exceed that for PG3. This is because goal count suffers from the same limitation as policy evaluation, to a lesser degree: all policies receive trivial scores until one is found that reaches at least one goal atom. Plan comparison has consistently good performance only in Ferry; in this domain alone, the plans generated by the planner are consistent with the compact policy that is ultimately learned. GNN BC is similarly constrained by the plans found by the planner. Combo improves only marginally on policy evaluation and plan comparison individually.

6 Conclusion

In this work, we proposed PG3 as a new approach for generalized planning. We demonstrated theoretically and empirically that PG3 outperforms alternative formulations of GPS, such as policy execution and plan comparison, and found that it is able to efficiently discover compact policies in PDDL domains. There are many interesting directions for future work, including applying PG3 in domains with stochasticity and partial observability, and integrating insights from other approaches to generalized planning.

References

- Bonet, B.; and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence*, 2001 Jun; 129 (1-2): 5-33.
- Bonet, B.; and Geffner, H. 2015. Policies that generalize: Solving many planning problems with the same policy. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Bonet, B.; and Geffner, H. 2018. Features, Projections, and Representation Change for Generalized Planning. *CoRR*, abs/1801.10055.
- Fikes, R. E.; Hart, P. E.; and Nilsson, N. J. 1972. Learning and executing generalized robot plans. *Artificial intelligence*, 3: 251–288.
- Gehring, C.; Asai, M.; Chitnis, R.; Silver, T.; Kaelbling, L.; Sohrabi, S.; and Katz, M. 2021. Reinforcement Learning for Classical Planning: Viewing Heuristics as Dense Reward Generators. In *International Conference on Automated Planning and Scheduling (ICAPS) PRL Workshop*.
- Groshev, E.; Goldstein, M.; Tamar, A.; Srivastava, S.; and Abbeel, P. 2018. Learning generalized reactive policies using deep neural networks. In *Proceedings of the International Conference on Automated Planning and Scheduling*.
- Hu, Y.; and Levesque, H. 2009. Planning with loops: Some new results. In *ICAPS Workshop on Generalized Planning*, 37.
- Illanes, L.; and McIlraith, S. A. 2019. Generalized Planning via Abstraction: Arbitrary Numbers of Objects. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 7610–7618.
- Jiménez, S.; and Jonsson, A. 2015. Computing plans with control flow and procedures using a classical planner. In *Proceedings of the Eighth Annual Symposium on Combinatorial Search, SOCS-15*, 62–69.
- Jiménez, S.; Segovia-Aguas, J.; and Jonsson, A. 2019. A review of generalized planning. *The Knowledge Engineering Review*, 34: e5.
- Karia, R.; and Srivastava, S. 2020. Learning generalized relational heuristic networks for model-agnostic planning. *arXiv preprint arXiv:2007.06702*.
- Khardon, R. 1999. Learning action strategies for planning domains. *Artificial Intelligence*, 113(1-2): 125–148.
- Levesque, H. 2005. Planning with Loops. In *IJCAI*.
- Levine, J.; and Humphreys, D. 2003. Learning action strategies for planning domains using genetic programming. In *Workshops on Applications of Evolutionary Computation*, 684–695. Springer.
- Martin, M.; and Geffner, H. 2004. Learning generalized policies from planning examples using concept languages. *Applied Intelligence*, 20(1): 9–19.
- McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. PDDL—the planning domain definition language.
- Mooney, R. J.; and Califf, M. E. 1995. Induction of first-order decision lists: Results on learning the past tense of English verbs. *Journal of Artificial Intelligence Research*, 3: 1–24.
- Pong, V.; Gu, S.; Dalal, M.; and Levine, S. 2018. Temporal difference models: Model-free deep rl for model-based control. *arXiv preprint arXiv:1802.09081*.
- Rivlin, O.; Hazan, T.; and Karpas, E. 2020. Generalized Planning With Deep Reinforcement Learning. *arXiv preprint arXiv:2005.02305*.
- Schaul, T.; Horgan, D.; Gregor, K.; and Silver, D. 2015. Universal value function approximators. In *International conference on machine learning*, 1312–1320. PMLR.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2018. Computing hierarchical finite state controllers with classical planning. *Journal of Artificial Intelligence Research*, 62: 755–797.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2021. Generalized Planning as Heuristic Search. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 31, 569–577.
- Silver, T.; and Chitnis, R. 2020. PDDL Gym: Gym Environments from PDDL Problems. In *International Conference on Automated Planning and Scheduling (ICAPS) PRL Workshop*.
- Silver, T.; Chitnis, R.; Curtis, A.; Tenenbaum, J.; Lozano-Perez, T.; and Kaelbling, L. P. 2021. Planning with Learned Object Importance in Large Problem Instances using Graph Neural Networks. In *AAAI*.
- Srivastava, S. 2011. Foundations and applications of generalized planning. *AI Communications*, 24(4): 349–351.
- Srivastava, S.; Immerman, N.; Zilberstein, S.; and Zhang, T. 2011a. Directed Search for Generalized Plans Using Classical Planners. In *ICAPS*.
- Srivastava, S.; Zilberstein, S.; Immerman, N.; and Geffner, H. 2011b. Qualitative Numeric Planning. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*, 1010–1016. San Francisco, California.
- Winner, E. Z. 2008. *Learning Domain-Specific Planners from Example Plans*. Ph.D. thesis, USA. AAI3345265.
- Yoon, S.; Fern, A.; and Givan, R. 2008. Learning control knowledge for forward search planning. *Journal of Machine Learning Research*, 9(Apr): 683–718.

A Proofs

Proof of Theorem 1

Theorem 1. *Under Assumption 1, in the tabular setting (Example 2), if a policy π solves all but one of the problems in Ψ , then the PG3 score for π is equal to the GPS cost-to-go.*

Proof. Let $\langle O, I, G \rangle$ be the problem in Ψ that π does not solve. First we will show that the PG3 score (Algorithm 2) is less than or equal to the GPS cost-to-go (Definition 1). Suppose that the PG3 score is n . Then for the problem $\langle O, I, G \rangle$, there exists a plan $(\underline{a}_1, \dots, \underline{a}_m)$ with corresponding states (S_1, \dots, S_{m+1}) , and indices $\mathcal{I} = \{i_1, \dots, i_n\}$ such that $\forall j \in \mathcal{I}, \pi(S_j) \neq \underline{a}_j$, and $\forall k \notin \mathcal{I}, \pi(S_k) = \underline{a}_k$. Consider the following tabular policy:

$$\{(S_j, G) \mapsto \underline{a}_j : j \in \mathcal{I}\} \cup \{(S, G) \mapsto \pi(S) : S \notin S_{\mathcal{I}}\}$$

where $S_{\mathcal{I}} = \{S_j : j \in \mathcal{I}\}$. By construction, this policy is satisficing. Furthermore, this policy differs from π in n table entries, so the cost-to-go is no more than n .

In the other direction, suppose towards a contradiction that the PG3 score is again n , but the GPS cost-to-go is $n' < n$. Observe that the successor function used in policy-guided planning induces a modified search graph with respect to the search graph in the original problem. Whereas all edge costs in the original search graph are unitary by Assumption 1, the modified search graph contains new zero-cost edges corresponding to the policy’s actions. Since the GPS cost-to-go is n' , there exists a satisficing policy π' that differs from π in n' states. Let $(\underline{a}_1, \dots, \underline{a}_m)$ be the plan for policy π' with corresponding states (S_1, \dots, S_{m+1}) . Since the policy is deterministic, these states do not have repeats (otherwise the policy could not be satisficing). The cost of this plan in the modified search graph is therefore equal to n' . However, by supposition, the minimum cost path in the modified search graph is n , a contradiction. Therefore we can conclude that the PG3 score is equal to the cost-to-go. \square

Proof of Theorem 2

Theorem 2. *Under Assumption 1, in the tabular setting (Example 2), PG3 is a lower bound on the GPS cost-to-go.*

Proof. Suppose to the contrary that PG3 is not GPS admissible. Then there exists a policy π and training problems Ψ where the PG3 score n exceeds the GPS cost-to-go n' . Since PG3 accumulates per-plan scores using a *max* (Algorithm 2), there is a corresponding problem $\langle O, I, G \rangle$ in Ψ that represents the *argmax*. In other words, the PG3 score for $\{\langle O, I, G \rangle\}$ is n . From Theorem 1, we know that the GPS cost-to-go for π on $\{\langle O, I, G \rangle\}$ is equal to n . Furthermore, the GPS cost-to-go n' for the full set Ψ cannot be less than n , since each additional problem contributes non-negatively to the GPS cost-to-go in the tabular setting. Thus $n \leq n'$, a contradiction. \square

B Induce Rule from Plans Operator

Here we describe in detail the final search operator, with a detailed walkthrough in the next subsection. This search operator proposes a modification to a given policy based on

example plans. The plans are cached from the score function when possible; for example, PG3 uses the policy-guided plans, and plan comparison uses the planner-generated plans. This operator suggests a modification to the candidate policy by identifying the last *missed action* — an action that does not match the candidate’s output at the corresponding *missed state* in one of the plans — and computing a new rule for the policy that would “correct” the policy on that example.

New rules are computed via an extension of the *triangle tables* approach of (Fikes, Hart, and Nilsson 1972), which was originally proposed for macro action synthesis. Given a plan subsequence, triangle tables creates a macro by computing a *preimage* for (an edited subsequence of) the subsequence. The plan subsequence that we use as input to triangle tables is one that begins with the missed action and ends with the first subsequent goal literal that is permanently achieved, that is, added to the state and never removed in the remainder of the plan. The preimage returned by triangle tables comprises the preconditions $\text{PRE}(\rho')$ of the new rule ρ' ; the goal conditions $\text{GOAL}(\rho')$ consist of the single goal literal at the tail of the subsequence; and the action $\text{ACT}(\rho')$ is derived from the missed action. The new rule ρ' is inserted into π before the first existing rule that was applicable in the missed state; if none were applicable, the new rule is appended to the end.

A limitation of the approach described thus far is that the preconditions added to the new rule can be overly specific. We would prefer to introduce a more general rule, to promote generalization to new problems, so long as that rule still returns the missed action in the missed state. With that intuition in mind, we implement a final extension of the triangle tables approach. Instead of starting with the entire preimage as the precondition set for the rule, we start with a smaller set, computed as follows. All preconditions in the preimage which only involve objects in the failed action and the first permanently-achieved subsequent goal literal are included in proposed rule. If upon checking that the proposed rule does not produce the missed action, we iteratively step through each action in the plan subsequence used in the triangle table and add new preconditions to that rule, attempting to reproduce the missed action. The new preconditions that we add at each iteration are determined as follows: we add all preconditions involving objects from the previously proposed rule and the current action in the plan subsequence. If the missed action is never produced, the final proposed rule whose preconditions are equivalent to the entire preimage is returned.

In preliminary experiments, we found that this operator is very helpful in most domains, such as the Ferry example shown in the following section. But in the Forest domain, the operator can lead to overfitting and syntactically large policies. We therefore disable this operator for Forest; note that this change applies to all GPS methods, including PG3, policy evaluation, plan comparison, combo, and goal count.

Example

We now provide an example of the Induce Rule from Plans operator in the Ferry domain. Consider the policy in Figure


```

(:rule rule1
 :parameters (?x1 ?x2)
 :state-preconditions (and
  (car ?x1)
  (location ?x2)
  (on ?x1)
 )
 :goal-preconditions (and
  (at ?x1 ?x2)
 )
 :action (debark ?x1 ?x2)
 )

```

Figure 2: Example input to Induce Rule from Plans operator; see Section B

```

(:rule rule1
 :parameters (?x1 ?x2)
 :state-preconditions (and
  (car ?x1)
  (location ?x2)
  (on ?x1)
 )
 :goal-preconditions (and
  (at ?x1 ?x2)
 )
 :action (debark ?x1 ?x2)
 )
(:rule rule2
 :parameters (?x1 ?x2 ?x3)
 :state-preconditions (and
  (car ?x1)
  (location ?x2)
  (location ?x3)
  (notEq ?x2 ?x3)
 )
 :goal-preconditions (and
  (not (at ?x1 ?x2))
  (atFerry ?x3)
  (on ?x1)
 )
 :action (sail ?x3 ?x2)
 )

```

Figure 3: Example output from Induce Rule from Plans operator; see Section B and Figure 2.

2 as input to the operator, and consider the following plan:

1. sail(18, 17)
2. board(c0, 17)
3. sail(17, 10)
4. debark(c0, 10)
5. sail(10, 12)
6. board(c4, 12)
7. sail(12, 18)
8. debark(c4, 18)

for a problem where $G = \{at(c4, 18), at(c0, 10)\}$. The initial state is too large to print, but includes `atFerry(18)`, `at(c0, 17)`, `at(c4, 12)` among others.

Referring to the policy, the last missed action in this plan is `sail(12, 18)`. The subsequent goal literal that is permanently achieved is `at(c4, 18)`, and it is achieved on the final time step, so the input to triangle tables is:

7. sail(12, 18)
8. debark(c4, 18)

The preimage then returned by triangle tables is:

```

{(car c4), (location 12), (location 18),
 (notEq 12, 18), (not (at 12, 18)), (atFerry
 (on c4))

```

Replacing all objects with placeholders, we arrive at a new rule, which we then insert into the policy following the existing rule; see Figure 3. In this example, the additional generalization step described in the section above does not change the policy, so the policy in Figure 3 is the final output of the operator.

C Additional GPS Optimizations

There are two additional optimizations that we use in experiments to accelerate GPS. First, in addition to the score functions described previously, we add a small *size penalty* to the score:

$$\text{PENALTY}(\pi) = -w \sum_{\rho \in \pi} |\text{PRE}(\rho)| + |\text{GOAL}(\rho)|$$

where w is a weight selected to be small enough (0.00001 in experiments) so that the penalty only breaks ties in the original score function. This regularization encourages GPS to

discover policies with shorter description length, encoding an Occam’s razor bias.

The second optimization is motivated by the observation that two policies may be syntactically distinct, but semantically equivalent with respect to all states and goals of interest. This leads to two challenges: (1) it is difficult to compute semantic equivalence, especially with respect to the unknown set of states and goals of interest; and (2) we cannot necessarily prune policies according to semantic equivalence because their neighbors in policy search space may be semantically different. Nonetheless, we would like to bias policy search so that it avoids redundancy.

To address the challenge of computing semantic equivalence, we resort to the following approximate approach. At the beginning of GPS, we use a planner to collect a set of example plans on the training problems (as in the plan comparison scoring function). The states and goals encountered in these example plans are used to approximate the “states and goals of interest.” During GPS, we execute each candidate policy on each of these states and goals to compute a *semantic identifier*. We maintain a map from semantic identifiers to counts, with the counts representing the number of times that a policy with that semantic identifier has been expanded. When a new policy is generated, it is inserted into the GBFS queue (Algorithm 1) with priority (count, score), with the effect being that semantically unique policies are always explored first. This optimization maintains the completeness of GPS while favoring policies that are approximately semantically unique.

D Additional Experimental Details

All code is implemented in Python and run on a quad-core AMD64 processor with 16GB RAM, in Ubuntu 18.04. We use PDDLgym (Silver and Chitnis 2020) for simulation.

Below are additional details for each domain and train/test problem distributions, which are procedurally generated separately for each random seed. The specific problems used in this work are supplied with the code.

- **Forest**

- *Max horizon*: 100

Domains	PG3 (Ours)		Policy Eval		Plan Comp		Combo		Goal Count		GNN BC		Random	
	Eval	Time	Eval	Time	Eval	Time	Eval	Time	Eval	Time	Eval	Time	Eval	Time
Delivery	0.00	0.8	0.00	–	0.32	–	0.42	1659.3	0.00	19.4	0.41	27.8	0.00	–
Forest	0.00	80.0	0.00	472.2	0.30	–	0.00	689.2	0.33	517.6	0.07	–	0.03	–
Gripper	0.00	20.8	0.00	–	0.42	0.2	0.00	–	0.00	–	0.13	–	0.00	–
Miconic	0.00	606.1	0.00	–	0.30	–	0.00	–	0.31	1810.6	0.12	–	0.05	–
Ferry	0.00	5.2	0.00	–	0.32	0.3	0.00	–	0.00	2340.0	0.18	–	0.00	–
Spanner	0.00	15.9	0.00	3.0	0.53	0.4	0.00	17.6	0.00	2.0	0.42	12.5	0.04	–

Table 2: **Standard deviations for policy learning results.** See Table 1 for further description.

- *Train problems*: 10 problems, which have grids with heights and widths ranging from 8 to 10, for a total of **64–100 objects**.
- *Test problems*: 30 problems, which have grids with heights and widths ranging from 10 to 12, for a total of **100–144 objects**.
- **Delivery**
 - *Max horizon*: 1000
 - *Train problems*: 5 problems, which have 5–9 locations, 2–4 delivery requests, and 2–4 packages, for a total of **9–17 objects**.
 - *Test problems*: 30 problems, which have 30–39 locations, 20–29 requests, and 0–10 more packages than requests, for a total of **70–110 objects**.
- **Miconic**
 - *Max horizon*: 1000
 - *Train problems*: 10 problems, which have 1–2 buildings, and for each building, 5–10 floors and 1–5 passengers, for a total of **6–30 objects**.
 - *Test problems*: 30 problems, which have 1–5 buildings, and for each building, 10–20 floors and 1–10 passengers, for a total of **11–150 objects**.
- **Gripper**
 - *Max horizon*: 1000
 - *Train problems*: 10 problems, which have 5–10 balls and 15–20 rooms for a total of **20–30 objects**.
 - *Test problems*: 30 problems, which have 20–30 balls and 40–50 rooms for a total of **60–80 objects**.
- **Ferry**
 - *Max horizon*: 1000
 - *Train problems*: 10 problems, which have 10–15 locations and 3–5 cars, for a total of **13–20 objects**.
 - *Test problems*: 30 problems, which have 20–30 locations and 10–20 cars, for a total **30–50 objects**.
- **Spanner**
 - *Max horizon*: 1000
 - *Train problems*: 10 problems, which have 3–5 spanners, 3–5 nuts, and 3–5 locations, for a total of **9–15 objects**.
 - *Test problems*: 30 problems, which have 10–20 spanners, 10–20 nuts, and 10–20 locations, for a total of **30–60 objects**.

E Additional Results

Standard Deviations for Main Results

See Table 2 for standard deviations on the main policy learning results presented in Table 1.

Qualitative Results: Learned Policies

See Figures 4, 5, 6, 7, 8, 9 for policies learned by PG3.

F Graph Neural Network Baseline Details

The GNN behavioral cloning baseline is implemented in PyTorch, version 1.9.0. Our architecture and learning hyperparameters are identical to those used in the reactive policy baseline of (Silver et al. 2021). All GNNs node and edge modules are fully connected neural networks with one hidden layer of dimension 16, ReLU activations, and layer normalization. Message passing is performed for $K = 3$ iterations. Training uses the Adam optimizer with learning rate 0.001 for 100000 epochs with a batch size of 16. Performance is evaluated by directly executing the learned policy on the test problems for up to 1000 steps. The reported validation accuracy is the percent of test problem executions that reach the goal after the last training epoch.

```

Forest Learned Policy

(:rule rule1
:parameters (?x1 - loc
             ?x2 - loc)
:state-preconditions (and
  (at ?x1)
  (not (isRock ?x2))
  (not (isWater ?x1))
  (adjacent ?x1 ?x2)
  (onTrail ?x1 ?x2)
)
:goal-preconditions ()
:action (walk ?x1 ?x2)
)

(:rule rule2
:parameters (?x1 - loc
             ?x2 - loc)
:state-preconditions (and
  (at ?x1)
  (isRock ?x2)
  (not (isWater ?x1))
  (adjacent ?x1 ?x2)
  (onTrail ?x1 ?x2)
)
:goal-preconditions ()
:action (climb ?x1 ?x2)
)

```

Figure 4: **Learned policy for the Forest domain.** This policy has the agent follow the marked trail until a rock is encountered, at which point it climbs over the rock. For this policy and the others, it is important to remember that the rules are *ordered* (left to right); each rule will only be used when the previous rules are not applicable.

```

Delivery Learned Policy

(:rule rule1
:parameters (?x1 - loc
             ?x2 - package)
:state-preconditions (and
  (not (satisfied ?x1))
  (at ?x1)
  (carrying ?x2)
)
:goal-preconditions ()
:action (deliver ?x2 ?x1)
)

(:rule rule2
:parameters (?x1 - loc
             ?x2 - package
             ?x3 - loc)
:state-preconditions (and
  (at ?x1)
  (isHomeBase ?x1)
  (safe ?x1)
  (unpacked ?x2)
)
:goal-preconditions (and
  (satisfied ?x3)
)
:action (pickup ?x2 ?x1)
)

(:rule rule3
:parameters (?x1 - loc
             ?x2 - loc)
:state-preconditions (and
  (at ?x1)
  (not (satisfied ?x2))
  (safe ?x1)
)
:goal-preconditions (and
  (satisfied ?x2)
)
:action (move ?x1 ?x2)
)

```

Figure 5: **Learned policy for the Delivery domain.** This policy has the agent deliver a package whenever it is at a location that wants a package but does not yet have one; it also picks up packages whenever there are packages available at the home base; and it moves to a location that wants a package, but has not yet received one.

Miconic Learned Policy

```

(rule rule1
:parameters (?x1 - floor
             ?x2 - passenger)
:state-preconditions (and
  (boarded ?x2)
  (destin ?x2 ?x1)
  (liftAt ?x1)
)
:goal-preconditions (and
  (served ?x2)
)
:action (depart ?x1 ?x2)
)

(rule rule2
:parameters (?x1 - floor
             ?x2 - passenger
             ?x3 - floor)
:state-preconditions (and
  (not (served ?x2))
  (boarded ?x2)
  (above ?x1 ?x3)
  (destin ?x2 ?x3)
  (liftAt ?x1)
)
:goal-preconditions (and
  (served ?x2)
)
:action (up ?x1 ?x3)
)

(rule rule3
:parameters (?x1 - floor
             ?x2 - passenger
             ?x3 - floor)
:state-preconditions (and
  (not (served ?x2))
  (boarded ?x2)
  (above ?x3 ?x1)
  (liftAt ?x1)
)
:goal-preconditions (and
  (served ?x2)
)
:action (down ?x1 ?x3)
)

(rule rule4
:parameters (?x1 - floor
             ?x2 - passenger)
:state-preconditions (and
  (not (served ?x2))
  (origin ?x2 ?x1)
  (liftAt ?x1)
)
:goal-preconditions (and
  (served ?x2)
)
:action (board ?x1 ?x2)
)

(rule rule5
:parameters (?x1 - floor
             ?x2 - passenger
             ?x3 - floor)
:state-preconditions (and
  (not (served ?x2))
  (above ?x1 ?x3)
  (origin ?x2 ?x3)
  (liftAt ?x1)
)
:goal-preconditions (and
  (served ?x2)
)
:action (up ?x1 ?x3)
)

```

Figure 6: **Learned policy for the Miconic domain.** This policy has the agent drop off passengers whenever they are boarded and at their destination floor; move up or down, depending on whether the destination floor is above or below, whenever a passenger is boarded but not yet at their destination floor; board a passenger whenever there is a passenger at the elevator's current floor that is not yet at their destination; and to start, moves up to a floor where there is a passenger waiting (the elevator always starts on the ground floor).

Gripper Learned Policy

```

(rule rule1
:parameters (?x1 ?x2 ?x3)
:state-preconditions (and
  (ball ?x1)
  (gripper ?x2)
  (room ?x3)
  (not (at ?x1 ?x3))
  (atRobby ?x3)
  (carry ?x1 ?x2)
)
:goal-preconditions (and
  (at ?x1 ?x3)
)
:action (drop ?x1 ?x3 ?x2)
)

(rule rule2
:parameters (?x1 ?x2 ?x3 ?x4)
:state-preconditions (and
  (ball ?x1)
  (gripper ?x2)
  (room ?x3)
  (room ?x4)
  (at ?x1 ?x3)
  (not (at ?x1 ?x4))
  (atRobby ?x3)
  (free ?x2)
)
:goal-preconditions (and
  (at ?x1 ?x4)
)
:action (pick ?x1 ?x3 ?x2)
)

(rule rule3
:parameters (?x1 ?x2 ?x3 ?x4)
:state-preconditions (and
  (ball ?x1)
  (gripper ?x2)
  (room ?x3)
  (room ?x4)
  (not (at ?x1 ?x4))
  (carry ?x1 ?x2)
  (atRobby ?x3)
)
:goal-preconditions (and
  (at ?x1 ?x4)
)
:action (move ?x3 ?x4)
)

(rule rule4
:parameters (?x1 ?x2 ?x3 ?x4)
:state-preconditions (and
  (ball ?x1)
  (room ?x2)
  (room ?x3)
  (room ?x4)
  (not (room ?x1))
  (at ?x1 ?x3)
  (not (at ?x1 ?x4))
  (atRobby ?x2)
)
:goal-preconditions (and
  (at ?x1 ?x4)
)
:action (move ?x2 ?x3)
)

```

Figure 7: **Learned policy for the Gripper domain.** Note that objects and parameters in this domain are untyped. This policy has the agent drop a ball that it's carrying when in the target room for that ball; pick a ball when the gripper is free and the ball is not yet in its target room; move to a target room when carrying a ball whose target room is not the robot's current room; and initially, moves to a room with some ball that is not yet at its target room.

Ferry Learned Policy

```

(:rule rule1
:parameters (?x1 ?x2)
:state-preconditions (and
(car ?x1)
(location ?x2)
(on ?x1)
)
:goal-preconditions (and
(at ?x1 ?x2)
)
:action (debark ?x1 ?x2)
)

(:rule rule2
:parameters (?x1 ?x2 ?x3)
:state-preconditions (and
(car ?x1)
(location ?x2)
(location ?x3)
(notEq ?x2 ?x3)
(not (at ?x1 ?x2))
(atFerry ?x3)
(on ?x1)
)
:goal-preconditions (and
(at ?x1 ?x2)
)
:action (sail ?x3 ?x2)
)

(:rule rule3
:parameters (?x1 ?x2 ?x3)
:state-preconditions (and
(car ?x1)
(location ?x2)
(location ?x3)
(notEq ?x2 ?x3)
(not (at ?x1 ?x2))
(atFerry ?x3)
(at ?x1 ?x3)
(empty-ferry)
)
:goal-preconditions (and
(at ?x1 ?x2)
)
:action (board ?x1 ?x2)
)

(:rule rule4
:parameters (?x1 ?x2 ?x3 ?x4)
:state-preconditions (and
(car ?x1)
(location ?x2)
(location ?x3)
(location ?x4)
(notEq ?x2 ?x3)
(notEq ?x3 ?x4)
(not (at ?x1 ?x2))
(atFerry ?x3)
(at ?x1 ?x4)
(empty-ferry)
)
:goal-preconditions (and
(at ?x1 ?x2)
)
:action (sail ?x3 ?x4)
)

```

Figure 8: **Learned policy for the Ferry domain.** Note that objects and parameters in this domain are untyped. This policy has the agent drop off (debark) cars whenever they are on the ferry but at their destinations; sail to a destination location when there is a car loaded; board a car when the ferry is empty and there is a car at the ferry's current location that has not yet been served; and sail to a location with a car that has not yet been served otherwise.

Spanner Learned Policy

```

(:rule rule1
:parameters (?x1 - person
?x2 - loc
?x3 - nut
?x4 - spanner)
:state-preconditions (and
(not (tightened ?x3))
(loose ?x3)
(at ?x1 ?x2)
(at ?x3 ?x2)
(carrying ?x1 ?x4)
(useable ?x4)
)
:goal-preconditions (and
(tightened ?x3)
)
:action (tighten ?x2 ?x4 ?x1 ?x3)
)

(:rule rule2
:parameters (?x1 - person
?x2 - loc
?x3 - nut
?x4 - spanner)
:state-preconditions (and
(not (tightened ?x3))
(loose ?x3)
(at ?x1 ?x2)
(at ?x4 ?x2)
(useable ?x4)
)
:goal-preconditions (and
(tightened ?x3)
)
:action (pickup ?x2 ?x4 ?x1)
)

(:rule rule3
:parameters (?x1 - person
?x2 - loc
?x3 - nut
?x4 - loc)
:state-preconditions (and
(not (tightened ?x3))
(loose ?x3)
(at ?x1 ?x2)
(at ?x3 ?x4)
(link ?x2 ?x4)
)
:goal-preconditions (and
(tightened ?x3)
)
:action (walk ?x2 ?x4 ?x1)
)

```

Figure 9: **Learned policy for the Spanner domain.** This policy has the agent tighten a nut whenever there is a loose nut at the agent's location and the agent is carrying a spanner (wrench); pickup a spanner whenever it is available at the current location and not yet carried; and walk to a location with a loose nut otherwise.