

---

# Learning Generalized Policy Automata for Relational Stochastic Shortest Path Problems

---

Rushang Karia, Rashmeet Kaur Nayyar, Siddharth Srivastava

School of Computing and Augmented Intelligence

Arizona State University

Tempe, AZ 85281

U.S.A.

{Rushang.Karia,rmnyayar,siddharths}@asu.edu

## Abstract

Several goal-oriented problems in the real-world can be naturally expressed as Stochastic Shortest Path problems (SSPs). However, the computational complexity of solving SSPs makes finding solutions to even moderately sized problems intractable. State-of-the-art SSP solvers are unable to learn generalized solutions or policies that would solve multiple problem instances with different object names and/or quantities. This paper presents an approach for learning *Generalized Policy Automata* (GPA): non-deterministic partial policies that can be used to catalyze the solution process. GPAs are learned using relational, feature-based abstractions, which makes them applicable on broad classes of related problems with different object names and quantities. Theoretical analysis of this approach shows that it guarantees completeness and hierarchical optimality. Empirical analysis shows that this approach effectively learns broadly applicable policy knowledge in a few-shot fashion and significantly outperforms state-of-the-art SSP solvers on test problems whose object counts are far greater than those used during training.

## 1 Introduction

Goal-oriented Markov Decision Processes (MDPs) expressed as Stochastic Shortest Path problems (SSPs) have been the subject of active research since they provide a convenient framework for modeling the uncertainty in action execution that often arises in the real-world. Recently, research in deep learning has demonstrated success in solving goal-oriented MDPs using image-based state representations (Tamar et al., 2016; Groshev et al., 2018). However, such methods require significant human-engineering effort in finding transformations like grayscale conversion, etc., to yield representations that facilitate learning. Many practical problems however, are more intuitively expressed using relational representations and have been widely studied in the literature.

As an example, consider a planetary rover whose mission is to collect all rocks of interest from a planet’s surface and deliver them to the base for analysis. Such a problem objective is not easily described in an image-based representation (e.g., visibility is affected by line of sight) but can be easily described using a relational description language such as first-order logic. Finding suitable image-based representations for such problems would be counter-productive and difficult. Furthermore, image-based deep learning methods often require large amounts of training data and/or are unable to provide guarantees of completeness and/or convergence.

Many real-world problems such as the rover example above can be readily expressed as SSPs using symbolic descriptions that can be solved in polynomial time in terms of the state space. SSP algorithms use a combination of pruning strategies (e.g., heuristics (Hansen and Zilberstein, 2001)) that can eliminate large parts of the search space from consideration thereby reducing the computational effort

expended. In spite of such optimizations, a major hurdle is the “curse-of-dimensionality” since the state spaces grow exponentially as the total number of objects increases. Existing SSP solvers would have difficulty scaling to rover problems with many locations and/or rocks. The pruning strategies employed by these SSP solvers do not scale well because they do not use knowledge that could have been exploited from solving similar problems. One solution to this problem is to compute a simple *generalized policy*: move the rover to the closest available location with an interesting rock, try loading the rock until it succeeds, navigate back to the base, unload it, and reiterate this process until all the interesting rocks are at the base. This generalized policy can be used to solve any rover instance with larger numbers of locations and rocks sharing a similar goal objective.

Related work in Generalized Planning addresses the problem of computing generalized policies by learning reliable controllers for broad classes of problems (Srivastava et al., 2008; Bonet et al., 2009; Aguas et al., 2016). Recently, Deep Learning based approaches have been shown to successfully learn generalized policies (Toyer et al., 2018; Garg et al., 2020; Karia and Srivastava, 2022). A key limitation of these approaches is the lack of any theoretical guarantees of finding a solution or optimality. In this paper, we show that such policies can be learned with guarantees of completeness and hierarchical optimality using solutions of very few, *small* problems with few objects.

The primary contribution of this paper is a novel approach for few-shot learning of *Generalized Policy Automata* (GPAs) using solutions of SSP instances with small object counts. GPAs are non-deterministic partial policies that represent generalized knowledge that can be applied to problems with different object names and larger object counts. This process uses logical feature-based abstractions to lift instance-specific information like object names and counts while preserving the relationships between objects in a way that can be used to express generalized knowledge. GPAs learned using our approach can be used to *accelerate any model-based SSP solver* by pruning out large sets of actions in different, related, but larger SSPs with many objects. We prove that our approach is complete and guarantees hierarchical optimality. Empirical analysis on a range of well-known benchmark domains shows that our approach few-shot learns GPAs using as few as three training problem instances and convincingly outperforms existing state-of-the-art SSP solvers and does so without compromising the quality of the solutions found.

The rest of this paper is organized as follows: The next section provides the necessary background. Sec. 3 describes our approach for using example policies in conjunction with abstractions to learn GPAs and using them for solving SSPs. We present our experimental setup and discuss obtained results in Sec. 4. Sec. 5 provides a description of related work in the area. Finally, Sec. 6 states the conclusions that we draw upon from this work followed by a brief description of future work.

## 2 Background

Our problem setting considers SSPs expressed in a symbolic description language such as the Probabilistic Planning Domain Definition Language (PPDDL) (Younes et al., 2005). Let  $\mathcal{D} = \langle \mathcal{P}, \mathcal{A} \rangle$  be a problem *domain* where  $\mathcal{P}$  and  $\mathcal{A}$  are finite sets of predicates and parameterized actions. Object *types*, such as those used in PPDDL, can be equivalently represented using unary predicates. A relational SSP problem instance for a domain  $\mathcal{D}$  with a goal formula  $g$  over  $\mathcal{P}$  and a finite set of objects  $O$  is defined as a tuple  $P = \langle O, S, A, T, C, s_0, g \rangle$ . A fact is the instantiation of a predicate  $p \in \mathcal{P}$  with the appropriate number of objects from  $O$ . A state  $s$  is a set of true facts and the state space  $S$  is defined as all possible sets of true facts derived using  $\mathcal{D}$  and  $O$ . Similarly, the action space  $A$  is instantiated using  $\mathcal{A}$  and  $O$ .  $T : S \times A \times S \rightarrow [0, 1]$  is the transition function and  $C : S \times A \times S \rightarrow \mathbb{R}^+$  is the cost function. An entry  $t(s, a, s') \in T$  defines the probability of executing action  $a \in A$  in a state  $s \in S$  and ending up in a state  $s' \in S$ , and  $c(s, a, s') \in C$  indicates the cost incurred while doing so. Naturally,  $\sum_{s'} t(s, a, s') = 1$  for any  $s \in S$  and  $a \in A$ . Note that  $a$  refers to the instantiated action  $a(o_1, \dots, o_n)$ , where  $o_1, \dots, o_n \in O$  are the action *parameters*. We omit the parameters when it is clear from context.  $s_0 \in S$  is a *known* initial state. A goal state  $s_g$  is a state such that  $s_g \models g$ . For any action  $a \in A$  and any goal state  $s_g \in S$ ,  $c(s_g, a, s_g) = 0$  and  $t(s_g, a, s_g) = 1$ . Additionally, termination (reaching a state  $s$  such that  $s \models g$ ) in an SSP is inevitable making the length of the horizon unknown but finite (Bertsekas and Tsitsiklis, 1996).

**Running example:** The planetary rover example can be expressed using a domain that consists of predicates  $connected(l_x, l_y)$ ,  $interesting(r_x)$ ,  $rover-at(l_x)$ ,  $in-rover(r_x)$ ,  $rock-at(r_x, l_x)$ , and parameterized actions  $load(r_x, l_x)$ ,  $unload(r_x, l_x)$ , and  $move(l_x, l_y)$ . Object types can be denoted using

unary predicates  $location(l_x)$  and  $rock(r_x)$ .  $l_x, l_y$ , and  $r_x$  are *parameters* that can be instantiated with different locations and rocks thus allowing for an easy way to express different problems. Actions dynamics are described using closed-form probability distributions (e.g. loading a rock could be modeled so that the rover picks up the rock with a probability of 0.8) and this forms the transition function. A simplified SSP problem that ignores connectivity and consists of two locations, a base location, and two rocks can be described using a set of objects  $O = \{l_1, l_2, l_{base}, r_1, r_2\}$ . A state in this SSP  $s_{eg}$  that describes the situation where  $r_2$  is being carried by the rover and  $r_1$  is at  $l_2$  can be written as  $s_{eg} = \{location(l_1), location(l_2), location(l_{base}), rock(r_1), rock(r_2), in-rover(r_2), rock-at(r_1, l_2)\}$ . The goal of delivering all the rocks to the base can be expressed as  $\forall x rock(x) \implies rock-at(x, l_{base})$ . Executing any action can be assumed to expend some fuel and as a result, the objective is to deliver all the rocks to the base in a way that minimizes the total fuel expended.

A solution to an SSP is a deterministic policy  $\pi : S \rightarrow A$ , that is a mapping from states to actions. A *proper policy* is one that is well-defined for all states. A *complete proper policy* is one for which termination is guaranteed from all possible states. By definition, SSPs must have at least one complete proper policy (Bertsekas and Tsitsiklis, 1996). This can be overly limiting in practice since such a formulation does not model dead end states: states from which the goal is reachable with probability 0. A weaker formulation of an SSP stipulates that the goal must be reachable with a probability of 1 from  $s_0$  i.e. whose solution is a *partial proper policy* that is well-defined for every state reachable from  $s_0$ . As such, we focus on a broader class of relaxed SSPs called Generalized SSPs (GSSPs) (Kolobov et al., 2012) that only require the existence of at least one partial proper policy. Henceforth, we use the term SSPs to refer to GSSPs and focus only on partial proper policies.

The value of a state  $s$  when using a policy  $\pi$  is the expected cost of executing  $\pi(s)$  when starting in  $s$  and following  $\pi$  thereafter:  $V^\pi(s) = \sum_{s' \in S} t(s, \pi(s), s')[c(s, \pi(s), s') + V^\pi(s')]$  (Sutton and Barto, 1998). Naturally,  $V^\pi(s_0) = \infty$  for any policy  $\pi$  that is not a partial proper policy.  $V$  is known as the value function. The optimal policy  $\pi^*$  is a policy that is better than or equal to all other policies.  $V^*$  is the optimal value function and can be described using the *Bellman optimality equations*:

$$V^*(s) = \min_{a \in A} \sum_{s' \in S} t(s, a, s')[c(s, a, s') + V^*(s')] \quad (1)$$

SSP solvers iteratively apply Eq. 1 starting from  $s_0$  to compute a policy, and under certain conditions, have been proved to converge to a policy that is  $\epsilon$ -consistent with  $\pi^*$  (Hansen and Zilberstein, 2001; Bonet and Geffner, 2003).

Let  $F_\alpha$  and  $F_\beta$  be two sets of features. We use feature-based abstractions to lift problem-specific characteristics like object names and numbers in order to facilitate the learning of generalized knowledge that can be applied to problems irrespective of differences in such characteristics. Given any SSP  $P$ , we define *state abstraction* as a function  $\alpha : F_\alpha, S \rightarrow \bar{S}$  that transforms the concrete state space  $S$  of  $P$  into a finite abstract state space  $\bar{S}$ . Similarly, *action abstraction*  $\beta : F_\beta, S, A \rightarrow \bar{A}$  transforms the action space to a finite abstract action space. Typically,  $|\bar{S}| \leq |S|$  and  $|\bar{A}| \leq |A|$ . We use  $\bar{s} = \alpha(F_\alpha, s)$  and  $\bar{a} = \beta(F_\beta, s, a)$  to represent abstractions of a concrete state  $s$  and action  $a$ . In this paper, we utilize feature sets automatically derived using canonical abstraction (Sagiv et al., 2002) to compute such feature-based representations of  $s$  and  $a$ . This is described in Sec. 3.1.

### 3 Our Approach

Our objective is to exploit knowledge from solutions of SSP instances with small object counts to learn Generalized Policy Automata (GPAs) that allow effective pruning of the search space of related SSPs with larger object counts. We accomplish this by using solutions to a small set of training instances that are easily solvable using existing SSP solvers, and using feature-based canonical abstractions to learn a GPA that encodes generalized partial policies and serves as a guide to prune the set of policies under consideration. We provide a brief description of canonical abstraction in Sec. 3.1, define GPAs in Sec. 3.2, and describe our process to learn a GPA in Sec. 3.2.1. We then describe our method (Alg. 1) for solving SSPs and state its theoretical properties in Sec. 3.3.

#### 3.1 Canonical Abstraction

Canonical abstractions, commonly used in program analysis, have been shown to be useful in generalized planning (Srivastava et al., 2011; Karia and Srivastava, 2021). Canonical abstractions

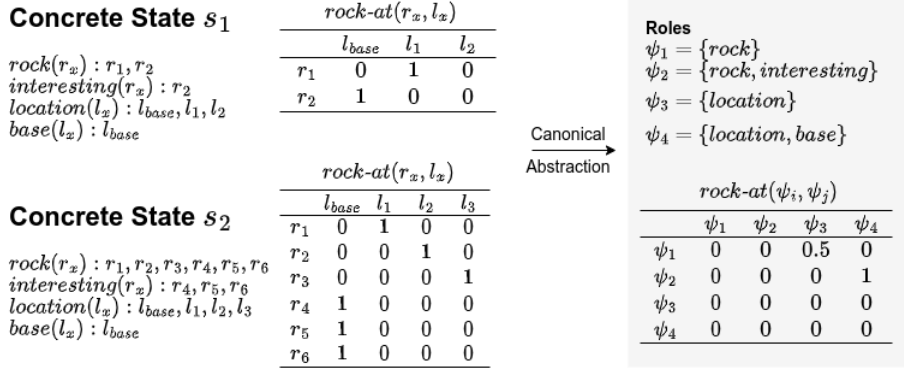


Figure 1: An example of how canonical abstraction can be used to lift problem-specific characteristics like object names and numbers.  $s_1$  and  $s_2$  are example states of two *different* problems.

group together multiple objects in a state using object roles. Given a concrete state  $s$  and an object  $o$ , the set of unary predicates that  $o$  satisfies is known as the *role* of  $o$ . 0-ary predicates are represented as unary predicates with a default ‘phantom’ object.

Let  $\psi$  be a role, then, we define  $\varphi_\psi(s)$  as a function that returns the set of objects that map to  $\psi$  in a concrete state  $s$ . Similarly, for any given predicate  $p_n \in \mathcal{P}$  where  $n$  is the arity,  $\varphi_{p_n(\psi_1, \dots, \psi_n)}(s)$  is defined as the set of all  $n$ -ary predicates in  $s$  that are consistent with the roles composing the predicate  $p_n(\psi_1, \dots, \psi_n)$ , i.e.,  $\varphi_{p_n(\psi_1, \dots, \psi_n)}(s) = \{p_n(o_1, \dots, o_n) \mid p_n(o_1, \dots, o_n) \in s, o_i \in \varphi_{\psi_i}(s)\}$ .

The value of a role  $\psi$  in a concrete state  $s$  is given as  $\max(2, |\varphi_\psi(s)|)$  to indicate whether there are 0, 1, or more than 1 objects satisfying  $\psi$ . Since relations between objects become imprecise when grouped as roles, the value of  $p_n(\psi_1, \dots, \psi_n)$  in  $s$  is determined using three-valued logic and is 0 if  $\varphi_{p_n(\psi_1, \dots, \psi_n)}(s) = \{\}$ , 1 if  $|\varphi_{p_n(\psi_1, \dots, \psi_n)}(s)| = |\varphi_{\psi_1}(s)| \times \dots \times |\varphi_{\psi_n}(s)|$ , and  $\frac{1}{2}$  otherwise.

Let  $\Psi$  be the set of all possible roles and  $\mathcal{P}_i$  be the set of predicate names  $p \in \mathcal{P}$  of arity  $i$  for a domain  $D$ , then  $\overline{\mathcal{P}}_i = \mathcal{P}_i \times [\Psi]^i$  is the set of all possible relations of arity  $i$  between roles. We define the feature set for state abstraction as  $F_\alpha = \Psi \cup_{i=2}^N \overline{\mathcal{P}}_i$  where  $N$  is the maximum arity of any predicate in  $D$ . We define state abstraction  $\alpha(F_\alpha, s)$  for a given concrete state  $s$  to return an abstract state  $\bar{s}$  as a total valuation of  $F_\alpha$  using the process described above. Similarly, we define the feature set for action abstraction as  $F_\beta = \Psi$ . The action abstraction  $\beta(F_\beta, s, a)$  for a concrete action  $a(o_1, \dots, o_n)$  when applied to  $s$  returns an abstract action  $\bar{a}(\psi_1, \dots, \psi_n)$  where the action name  $\bar{a} \equiv a$  and  $\psi_i$  is the role that the parameter  $o_i$  satisfies, i.e.,  $o_i \in \varphi_{\psi_i}(s)$  for any  $\psi_i \in \Psi$ .

Fig. 1 provides an intuitive example of how canonical abstraction can be used to lift instance-specific information. The figure describes two concrete states  $s_1$  and  $s_2$  from two different problems.  $s_1$  contains 2 rocks and 3 locations whereas  $s_2$  contains 6 rocks and 4 locations. There are four roles in  $\bar{s}$ . For example, in  $s_1$ ,  $r_1$  maps to the the role for a *rock*  $\psi_1$ , and  $r_2$  maps to the role for an *interesting rock*  $\psi_2$ , i.e.,  $\varphi_{\psi_1}(s_1) = \{r_1\}$  and  $\varphi_{\psi_2}(s_1) = \{r_2\}$ . The abstract relation  $rock-at(\psi_i, \psi_j)$  provides the three-valued representation of values between different roles. For example,  $rock-at(\psi_1, \psi_3)$  is interpreted as *the set of rocks that are at some location*, while  $rock-at(\psi_2, \psi_4)$  is interpreted as *the set of interesting rocks that are at the base*. Since  $rock-at(r_1, l_2)$  does not appear in  $s_1$ ,  $rock-at(\psi_1, \psi_3)$  evaluates to 0.5 to indicate that there are some *rocks* that are at some, but not all, *locations*. Similarly, since all *interesting rocks* are at the base location,  $rock-at(\psi_2, \psi_4)$  evaluates to 1.

The key aspect of abstraction comes from the observation that the relation  $rock-at(\psi_i, \psi_j)$  remains the same for  $s_2$  even though  $s_2$  has many more objects than  $s_1$ . The same high-level interpretations of the relations are captured while lifting low-level information like object names and numbers.

### 3.2 Generalized Policy Automata

We introduce Generalized Policy Automata (GPAs) as compact and expressive non-deterministic finite-state automata that encode generalized knowledge and can be represented as directed hypergraphs.

GPA's impose hierarchical constraints on the state space of an SSP and prune the action space under consideration, thus reducing the computational effort of solving larger related SSP instances.

**Definition 3.1** (Generalized Policy Automaton). Let  $\bar{S}$  and  $\bar{A}$  be a set of *abstract* states and actions. A Generalized Policy Automaton (GPA)  $\bar{\mathcal{G}} = \langle \bar{\mathcal{V}}, \bar{\mathcal{E}} \rangle$  is a non-deterministic finite-state automaton that can be represented as a directed hypergraph where the set of vertices  $\bar{\mathcal{V}} = \bar{S}$ .  $\bar{\mathcal{E}} \subseteq \bar{\mathcal{V}} \times \mathbb{P}(\bar{\mathcal{V}}) \setminus \emptyset \times \bar{A}$  where  $\mathbb{P}(\bar{\mathcal{V}})$  is the powerset of  $\bar{\mathcal{V}}$  is a set of directed hyperedges s.t. each hyperedge  $\bar{e} \in \bar{\mathcal{E}}$  is a tuple  $(\bar{e}_{src}, \bar{e}_{dest}, \bar{e}_{act})$  representing a start vertex, a set of result vertices, and an action label.

**Definition 3.2** (GPA Consistent Policy). A policy  $\pi$  for an SSP  $P$  is defined to be consistent with a GPA  $\bar{\mathcal{G}}$  iff for any states  $s, s' \in S$  and any action  $a \in A$  whenever  $\pi(s) = a$  and  $t(s, a, s') > 0$  there exists a hyperedge  $\bar{e} \equiv (e_{src}, e_{dest}, e_{act}) \in \bar{\mathcal{E}}$  where  $\bar{s} = \bar{e}_{src}$ ,  $\bar{s}' \in \bar{e}_{dest}$ , and  $\bar{a} = \bar{e}_{act}$ .

### 3.2.1 Learning GPAs

It is well-known that solutions to small problems can be used to construct generalized control structures that can assist in solving larger problems. We adopt a similar strategy of the learn-from-small-examples approach (Wu and Givan, 2007; Karia and Srivastava, 2021, 2022) and compute GPAs iteratively from a small training set containing solutions of similar SSP instances.

To form our training set, we use a library of solution policies  $\Pi = \{\pi_1, \dots, \pi_n\}$  for *small* SSPs  $P_1, \dots, P_n$  that can be easily (and optimally) computed by existing SSP solvers. We use the transition function for  $P_i$  to convert each policy  $\pi_i \in \Pi$  to a set of transitions  $\tau_i = \{(s, a, s') \mid s, s' \in S_i, a \in A_i, \pi_i(s) = a, t_i(s, a, s') > 0\}$ . We then construct our training set  $\mathcal{T} = \tau_1 \cup \dots \cup \tau_n$  as a set of concrete transitions.

The GPA  $\bar{\mathcal{G}}_\Pi$  learned from a finite set of concrete transitions  $\mathcal{T}$  is defined as follows. We first initialize a GPA  $\bar{\mathcal{G}}_\Pi = \langle \{\}, \{\} \rangle$ . Next, we convert  $\mathcal{T}$  into an abstract transition set  $\bar{\mathcal{T}} = \{(\alpha(F_\alpha, s), \beta(F_\beta, s, a), \alpha(F_\alpha, s')) \mid (s, a, s') \in \mathcal{T}\}$ . We then form the vertex set by using all abstract states in  $\bar{\mathcal{T}}$ , i.e.,  $\bar{\mathcal{V}} = \bar{\mathcal{V}} \cup \{\bar{s}, \bar{s}'\}$  for every  $(\bar{s}, \bar{a}, \bar{s}') \in \bar{\mathcal{T}}$ . Similarly, we convert each abstract transition into a hyperedge and add it to the edge set, i.e.  $\bar{\mathcal{E}} = \bar{\mathcal{E}} \cup (\bar{s}, \{\bar{s}'\}, \bar{a})$  for every  $(\bar{s}, \bar{a}, \bar{s}') \in \bar{\mathcal{T}}$ . Once this is done, we compress  $\bar{\mathcal{G}}_\Pi$  by replacing any edges in  $\bar{\mathcal{E}}$  that have the same start nodes and labels but different destinations with a single new edge that combines the destinations of the edges, i.e., for any two edges  $\bar{e}^1, \bar{e}^2 \in \bar{\mathcal{E}}$  s.t.  $\bar{e}^1_{src} = \bar{e}^2_{src}$ ,  $\bar{e}^1_{dest} \neq \bar{e}^2_{dest}$ , and  $\bar{e}^1_{act} = \bar{e}^2_{act}$ ,  $\bar{\mathcal{E}} = \bar{\mathcal{E}} \setminus \{\bar{e}^1, \bar{e}^2\} \cup (\bar{e}^1_{src}, \bar{e}^1_{dest} \cup \bar{e}^2_{dest}, \bar{e}^1_{act})$ . Fig. 2 provides a high-level overview of our procedure for merging GPAs. Henceforth, we drop the subscript  $\Pi$  from  $\bar{\mathcal{G}}_\Pi$  when it is clear from context.

### 3.3 Solving SSPs using GPAs

The key intuition behind our method is to use the GPA to prune policies that are not consistent with the GPA from the search process. We accomplish this by solving a GPA constrained problem that satisfies the constraints encoded by the GPA. We define a GPA constrained problem as follows:

**Definition 3.3** (GPA constrained problem). Let  $P = \langle O, S, A, T, C, s_0, g \rangle$  be an SSP for a domain  $D$  and let  $\bar{\mathcal{G}} = \langle \bar{\mathcal{V}}, \bar{\mathcal{E}} \rangle$  be a GPA. A GPA constrained problem  $P|_{\bar{\mathcal{G}}} = \langle O, S, A, T, C', s_0, g \rangle$  is defined with a cost function  $C' : S \times A \times S \rightarrow \mathbb{R}^+$  such that  $C'[s, a, s'] = C[s, a, s']$  when there exists a hyperedge  $\bar{e} \equiv (e_{src}, e_{dest}, e_{act}) \in \bar{\mathcal{E}}$  where  $\bar{s} = \bar{e}_{src}$ ,  $\bar{s}' \in \bar{e}_{dest}$ , and  $\bar{a} = \bar{e}_{act}$  and  $C'[s, a, s'] = \infty$  if there is no such hyperedge.

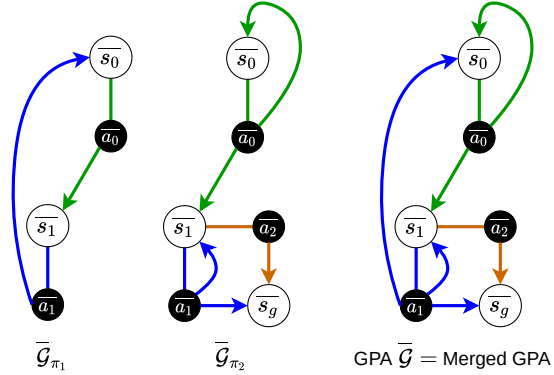


Figure 2: A high-level overview of how we merge different GPAs. All edges with the same color represent a hyperedge. For example, the blue colored hyperedge in  $\bar{\mathcal{G}}$  is  $(\bar{s}_1, \{\bar{s}_0, \bar{s}_g\}, \bar{a}_1)$ .

---

**Algorithm 1** GPA acceleration for SSPs

---

**Require:** SSP  $P = \langle O, S, A, T, C, s_0, g \rangle$ , GPA  $\bar{\mathcal{G}} = \langle \bar{\mathcal{V}}, \bar{\mathcal{E}} \rangle$ ,  
Feature Sets  $F_\alpha, F_\beta$ , Abstraction Functions  $\alpha, \beta$

- 1:  $C' = C$  {copy over the cost function of  $P$ }
- 2: **for**  $(s, a, s') \in S \times A \times S$  **do**
- 3:    $\bar{s}, \bar{a}, \bar{s}' \leftarrow \alpha(F_\alpha, s), \beta(F_\beta, s, a), \alpha(F_\alpha, s')$
- 4:   **if** there is no edge  $\bar{e} \in \bar{\mathcal{E}}$  s.t.  $\bar{e}_{src} = \bar{s}, \bar{s}' \in \bar{e}_{dest}$ , and  $\bar{e}_{act} = \bar{a}$  **then**
- 5:      $C'[s, a, s'] = \infty$
- 6:   **end if**
- 7: **end for**
- 8:  $P' = \langle O, S, A, T, C', s_0, g \rangle$
- 9:  $V_{P'}^* \leftarrow \text{initializeValueFunction}()$  {Typically using heuristics or randomly}
- 10:  $V_{P'}^*, \pi_{P'}^* \leftarrow \text{optimallySolveSSP}(P', V_{P'}^*)$
- 11: **if**  $\pi_{P'}^*$  is a partial proper policy **then**
- 12:   **return**  $\pi_{P'}^*$  {Return constrained policy}
- 13: **else**
- 14:    $V_P, \pi_P \leftarrow \text{optimallySolveSSP}(P, V_{P'}^*)$
- 15:   **return**  $\pi_P$  {Return partial proper policy for  $P$ }
- 16: **end if**

---

The goal of modifying the cost function is to prevent concrete transitions whose abstract translations are absent in the GPA from being used when performing Bellman updates for  $P|_{\bar{\mathcal{G}}}$ . Actions belonging to such transitions cannot appear in  $\pi_{P|_{\bar{\mathcal{G}}}}^*$  since their costs would be  $\infty$ . As a result,  $P|_{\bar{\mathcal{G}}}$  is not an SSP and the existence of a partial proper policy is not guaranteed in  $P|_{\bar{\mathcal{G}}}$ . Note that every optimal policy  $\pi_{P|_{\bar{\mathcal{G}}}}^*$  for  $P|_{\bar{\mathcal{G}}}$  is *hierarchically optimal* for an SSP  $P$  given a GPA  $\bar{\mathcal{G}}$ . Our overall objective is to compute such hierarchically optimal policies in a way such that they are approximately as cost effective as the optimal policy while requiring a fraction of the computational effort.

Given a GPA  $\bar{\mathcal{G}}$ , Alg. 1 works as follows. Lines 1-8 create a GPA constrained problem  $P'$ . Next, line 10 optimally solves  $P'$  using any off-the-shelf SSP solver with a randomly initialized value function (line 9). If the computed policy  $\pi_{P'}^*$  is a partial proper policy then it is returned immediately (lines 11-12) else Alg. 1 uses a new instance of the SSP solver to compute a policy for  $P$  using  $V_{P'}^*$  as the *bootstrapped* initial value function. Information such as whether a state is a dead end, etc., is not carried over. Line 15 then optimally solves  $P$  using  $V_{P'}^*$  as the initial value estimates and returns a partial proper policy  $\pi_P$  for  $P$ . Since  $V_{P'}^*$  is only used as an initial bootstrapping estimate for  $P$ , an SSP solver will only return a policy  $\pi_P$  that is better than or equal to  $\pi_{P'}^*$  following standard results on policy improvement for value iteration (Sutton and Barto, 1998). The following result shows that Alg. 1 computes hierarchically optimal policies.<sup>1</sup>

**Theorem 1.** When Alg. 1 returns a constrained policy  $\pi_{P'}^*$  (line 12) for an SSP  $P$  given a GPA  $\bar{\mathcal{G}}$  then  $\pi_{P'}^*$  is hierarchically optimal, i.e.,  $P' \equiv P|_{\bar{\mathcal{G}}}$ .

*Proof (Sketch).* Intuitively, Alg. 1 creates a new problem  $P'$  (lines 1-8) using the original cost function of  $P$ . The cost function for a concrete transition  $(s, a, s') \in S \times A \times S$  is set to  $\infty$  iff the corresponding abstract transition is not part of a hyperedge in  $\bar{\mathcal{G}}$ . As a result,  $P' \equiv P|_{\bar{\mathcal{G}}}$ , and any partial policy  $\pi_{P'}$  for  $\bar{\mathcal{G}}$  is a GPA consistent policy. Since Alg. 1 optimally solves  $P'$  (line 10), the returned policy is hierarchically optimal.  $\square$

Alg. 1 computes  $\pi_{P|_{\bar{\mathcal{G}}}}^*$  for  $P|_{\bar{\mathcal{G}}}$  in the space of cross-product of the states of the GPA  $\bar{\mathcal{G}}$  with the states of  $P$ , similar to HAMs (Parr and Russell, 1997). As seen in our empirical analysis in Sec. 4, we observe that a small set of example policies are sufficient to capture rich generalized control structures that are encoded by such hierarchically optimal policies. Alg. 1 computes such policies within a fraction of the original computational effort and in most cases with costs comparable to  $\pi_P^*$ .

The following result shows that Alg. 1 is complete in that it will always return a partial proper policy for an SSP (lines 12 and 14).

---

<sup>1</sup>Complete proofs are available in the extended version (Karia et al., 2022).

**Theorem 2.** Given a GPA  $\bar{\mathcal{G}}$  and an SSP  $P$ , Alg. 1 always returns a partial proper policy.

The following result indicates that the output of Alg. 1 is hierarchically optimal or better in terms of the expected value at  $s_0$ .

**Theorem 3.** Let  $V^\pi$  be the value function for a policy  $\pi$  returned by Alg. 1 for an SSP  $P$  using GPA  $\bar{\mathcal{G}}$ . Let  $V_{P|\bar{\mathcal{G}}}^*$  be the optimal value function for  $P|\bar{\mathcal{G}}$ , then  $V^\pi(s_0) \leq V_{P|\bar{\mathcal{G}}}^*(s_0)$ .

*Proof (Sketch).* If Alg. 1 finds a partial proper policy  $\pi_{P|\bar{\mathcal{G}}}^*$  for  $P|\bar{\mathcal{G}}$  then it returns it immediately (line 11) in which case  $V^\pi(s_0) = V_{P|\bar{\mathcal{G}}}^*(s_0)$ . If  $\pi_{P|\bar{\mathcal{G}}}^*$  is not a partial proper policy then  $V_{P|\bar{\mathcal{G}}}^*(s_0) = \infty$ . Since Alg. 1 is complete (Thm. 2),  $\pi$  is a partial proper policy where  $V^\pi(s_0) < \infty$ .  $\square$

**Corollary 3.1.** If  $V_P^*(s_0) = V_{P|\bar{\mathcal{G}}}^*(s_0)$ , then Alg. 1 returns the optimal policy for  $P$ .

The following result indicates that only a finite set of training examples are needed to learn a GPA such that the constrained problem will always yield the optimal policy for a given domain  $D$ .

**Theorem 4.** Suppose  $D$  is a domain,  $g$  is a formula over the predicates in  $D$ 's vocabulary,  $\bar{\mathcal{G}}_{\Pi^*}$  is a GPA s.t. for every SSP instance  $P$  of  $D$  whose goal is  $g$ , there exists an optimal policy  $\pi_P^*$  that is consistent with  $\bar{\mathcal{G}}_{\Pi^*}$ . Then there exists a finite set of policies  $\Pi^*$  from which  $\bar{\mathcal{G}}_{\Pi^*}$  can be learned.

*Proof (Sketch).* Since the sizes of the abstract state and action spaces  $\bar{\mathcal{S}}$  and  $\bar{\mathcal{A}}$  are finite, the size of  $\bar{\mathcal{G}}_{\Pi^*}$  is finite and is bounded by  $|\bar{\mathcal{S}}| \times |\bar{\mathcal{A}}| \times |\bar{\mathcal{S}}|$ . For every  $\bar{e} \in \bar{\mathcal{E}}$  at most  $|\bar{e}_{dest}|$  different transitions are needed to learn  $\bar{e}$ . Since there is a finite number of edges, a finite amount of training data will suffice for learning  $\bar{\mathcal{G}}_{\Pi^*}$ .  $\square$

In the worst case, the cost functions of  $P$  and  $P|\bar{\mathcal{G}}$  are similar and no savings are obtained. However, in our empirical evaluation, we observed that typically very few and small training problems suffice to learn a compact GPA  $\bar{\mathcal{G}}$  that allows efficient computation of solutions for problems that are significantly larger than those used during training. Finding the right set of examples  $\Pi^*$  from which  $\bar{\mathcal{G}}_{\Pi^*}$  can be learned is an interesting and non-trivial research problem that we leave to future work.

## 4 Experiments

We conducted an empirical evaluation on five well-known benchmark domains that were selected from the International Planning Competition (IPC) (Long and Fox, 2003), International Probabilistic Planning Competition (IPPC) (Younes et al., 2005), and robotic planning (Shah et al., 2020). As a part of our analysis, we use the time required to compute a solution and measure the quality of the solutions found to determine whether GPAs allow efficient solving of SSPs.

We chose PPDDL as our representational language, which was the default language in IPPCs until 2011, after which the Relational Dynamic Influence Diagram Language (RDDL) (Sanner, 2010) became the default. We chose PPDDL over RDDL since RDDL does not allow specifying the goal condition easily and as a result many benchmarks using RDDL are general purpose MDPs with no goals, and since modern state-of-the-art solvers for PPDDL are available.

For our baselines, we focus on complete solvers for SSPs. Deep Learning based approaches do not guarantee completeness and thus are not directly comparable with our work. We used Labeled RTDP (LRTDP) (Bonet and Geffner, 2003) and Soft-FLARES (Pineda and Zilberstein, 2019) which are state-of-the-art (SOA), complete SSP solvers. These algorithms internally generate their own heuristics for initializing the value function using the input domain and problem file. We used the inadmissible FF heuristic (Hoffmann, 2001) as the internal heuristic function for all algorithms since the baselines performed best using it.

We ran our experiments on a cluster of Intel Xeon E5-2680 v4 CPUs running at 2.4 GHz with 16 GiB of RAM. Our implementation is in Python and we ported C++ implementations of the baselines from `mdp-lib` (Pineda and Zilberstein, 2019) to Python.<sup>2</sup> We utilized problem generators from the IPC and IPPC suites and those in Shah et al. (2020) for generating the training and test problems for all domains. We provide a brief description of the problem domains below.

**Rover**( $r, w, s, o$ ) A set of  $r$  rovers need to collect and drop  $s$  samples that are present at one of  $w$

<sup>2</sup>Our source code is available at <https://github.com/AAIR-lab/GRAPL>




Figure 3: Impact of learned GPAs on solver performance (lower values better). Left y-axes and bars show solution times (in units of 1000 secs) for our approach and baseline SOA solvers (LRTDP and Soft-FLARES). Right y-axes and points show cost incurred by the policy computed by our approach and the baselines. We use the same SSP solver as the corresponding baseline in our approach. Error bars indicate 1 standard deviation (SD) averaged across 10 runs. For clarity, we only report results for the largest test problems and omit error bars from costs due to the low SDs. Open bars at the top indicate a timeout. A complete view of our empirical results is available in the appendix.

waypoints. The rovers also need to collect images of different objectives  $o$  that are visible from certain waypoints. This is an IPC domain and we converted it into a stochastic version by modifying sample collecting actions to fail with a probability of 0.4 (keeping the rover in the same state).

**Gripper**( $b$ ) A robot with two grippers is placed in an environment consisting of two rooms A and B. The objective of the robot is to transfer all the  $b$  balls initially located in room A to room B. We modified the gripper to be slippery so that picking a ball has a 20% chance of failure.

**Schedule**( $C, p$ ) is an IPPC domain that consists of a set of  $p$  packets, each belonging to one of  $C$  different classes that need to be queued. A router must first process the arrival of a packet in order to route it. The interval at which the router processes arrivals is determined by probability 0.94.

**Keva**( $P, h$ ) A robot uses  $P$  keva planks to build a tower of height  $h$ . Planks are placed in a specific order in one of two locations, preferring one location with probability 0.6. Despite this simple setting, the Keva domain has been shown to be a challenging problem in robotics (Shah et al., 2020).

**Delicate Can**( $c$ ) An arrangement of  $c$  cans on a table of which one is a delicate can. The objective is to pick up a specific goal can. Cans can obstruct the trajectory to the goal can and they must be moved in order to successfully pick up the goal can. Cans can be crushed with probability 0.1 (delicate cans have a higher chance with probability 0.8) and need to be revived.

**Training Setup** Our method learns GPAs in a few-shot fashion requiring little to no training data. For forming our training set, we used a minimum of three and a maximum of ten different solution policies (obtained using LAO\*) for each domain. The time required to learn a GPA was less than 10 seconds in all cases in our experiments. Our training speed highlights the advantages of GPAs that can be quickly learned in a few-shot setting. Moreover, compared to neuro-symbolic methods, GPAs are not subject to catastrophic forgetting and new training data can easily be merged with the existing GPA using the process described in Sec. 3.2.1.

**Test Setup** We fixed the time and memory limit for each problem to 7200 seconds and 16 GiB respectively. To demonstrate generalizability, our test set contains problems with object counts that are much larger than the training policies used. The largest problems in our test sets contain at least twice the number of objects than those used in training. For example, in the Keva domain, we use training policies with towers of height up to 6 and evaluate on problems with towers of height up to 14. The minimum and maximum number of problems that we used in our test set are 11 and 26 problems respectively. Due to space limitations, information pertaining to the total number of training and test problems and their parameters, hyperparameters used for configuring baselines, etc., is included in the appendix.

#### 4.1 Results and Analysis

Our evaluation metric compares the time required to find a partial proper policy. We also compare the quality of the computed policies by executing the policies for 100 trials with a horizon limit of



100 and averaging the obtained costs. We report our overall results averaged across 10 different runs and report results up to one standard deviation. Results of our experiments are illustrated in Fig. 3.

In four out of five domains (Schedule, Rover, Keva, and Delicate Can), our approach takes significantly less time compared to the corresponding baseline. For example, in Schedule, the baselines timed out for all of the large test problems reported. GPAs are able to successfully prune away action transitions that are not helping, leading to large savings in the computational effort. The costs obtained for executing these policies are also quite similar to each baseline (e.g. Keva), showing that GPAs are capable of learning good policies much faster without compromising solution quality.

Our approach was unable to outperform the baselines in the Gripper domain. An interesting phenomenon that we observed was that training policies returned by LAO\* were different for the case of even/odd balls due to tie-breaking and this led to the GPA not pruning actions as effectively. Nevertheless, we expected the GPA to outperform the baselines. We performed a deeper investigation and found that the FF heuristic used is already well-suited to prune away actions that the GPA would have otherwise pruned. This results in additional overhead being added in our SSP solver from the process of abstraction. However, heuristics that allow such pruning are difficult to synthesize, and in many cases, are hand-coded by an expert after employing significant effort.

There was a decrease in the gains observed in the Delicate Can domain between problem IDs 10 and 11. We investigated this issue and found that the GPA constructed was unable to effectively prune transitions in problem ID 11. This is because the training policies that were used were not sufficient to learn a GPA that could generalize well to larger problems in the Delicate Can domain. As a result,  $\pi_{P|G}^*$  was not a partial proper policy for problem ID 11 and thus additional computation was required.

Finally, because of the fixed timeout used, the maximum time of the baselines was bounded, making the impact of GPAs appear smaller in larger problems. For example, in problem ID 25 of the Rover domain, the Soft-FLARES baseline timed out in all our runs, but when allowed to run to convergence, it took over 15000 seconds in a targeted experiment that we performed for investigating this issue.

## 5 Related Work

There has been plenty of dedicated research to improve the efficiency for solving SSPs. LAO\* (Hansen and Zilberstein, 2001) computes policies by using heuristics to guide the search process. LRTDP (Bonet and Geffner, 2003) uses a labeling procedure in RTDP wherein a part of the subtree that is  $\epsilon$ -consistent is marked as *solved* leading to faster ending of trials. SSiPP (Trevizan and Veloso, 2012) uses short-sightedness by only considering reachable states up to  $t$  states away and solving this constrained SSP. Soft-FLARES (Pineda and Zilberstein, 2019) combines labeling and short-sightedness for computing solutions. These approaches are complete and can be configured to return optimal solutions, however, they do not learn any generalized knowledge and as result cannot readily scale to problems with a larger number of objects.

Determinization-based approaches (Yoon et al., 2007; Pineda and Zilberstein, 2014) build sparse representations of SSP problems by reducing the branching factor in the “environment’s choices” (the set of probabilistic effects of an action), while our approach uses abstraction to create abstract controllers that generalize solutions to SSPs and reduces the branching factor in the agent’s choice (the set of applicable actions). Our approach always considers all possible outcomes of every action. This is a key advantage of our approach since GPAs are able to better handle unexpected outcomes when executing an action in the policy that would otherwise require replanning.

Boutilier et al. (2001) utilize decision-theoretic regression to compute generalized policies for first-order MDPs represented using situation calculus. They utilize symbolic dynamic programming to compute a symbolic value function that applies to problems with varying number of objects. FOALP (Sanner and Boutilier, 2005) uses linear programming to compute an approximation of the value function for first-order MDPs while providing upper bounds on the approximation error irrespective of the domain size. A key limitation of their approach is requiring the use of a representation of action models over which it is possible to regress using situation calculus. API (Fern et al., 2006) uses approximate policy iteration with taxonomic decision lists to form policies. They use Monte Carlo simulations with random walks on a single problem to construct a policy. API offers no guarantees of completeness or hierarchical optimality.

Parr and Russell (1997) propose the hierarchical abstract machine (HAM) framework wherein component solutions from problem instances can be combined to solve larger problem instances efficiently. Recently, Bai and Russell (2017) extended HAMs to Reinforcement Learning settings by leveraging internal transitions of the HAMs. A key limitation of both these approaches is that the HAMs were hand-coded by a domain expert.

Related work in Generalized Planning focuses on the problem of computing generalized plans and policies such as our GPAs (Srivastava et al., 2012). Bonet et al. (2009) automatically create finite-state controllers for solving problems using a set of examples by modeling the search as a contingent problem. Their approach is limited in applicability since it only works on deterministic problems and the features they use are hand-coded. Aguas et al. (2016) utilize small example policies to synthesize hierarchical finite state controllers that can call each other. However, their approach requires all training data to be provided upfront. D2L (Francès et al., 2021) utilizes description logics to automatically generate features and reactive policies based on those features. Their approach comes with no guarantees for finding a solution and can only work on deterministic problems.

Deep Learning based approaches such as ASNs (Toyer et al., 2018) learn generalized policies for SSPs using a neuro-symbolic approach. They use the action schema from PPDDL to create alternating action and proposition layers. They do not learn generalized controllers and instead duplicate weights in a post-processing step to represent the generalized policy. GRL (Karia and Srivastava, 2022) uses Description Logic based feature evaluations of states as input to a neural network for computing generalized policies in settings where transition dynamics are not available as closed-form probability distributions. A common limitation of such approaches is the lack of any theoretical guarantees of completeness or optimality. Furthermore, these approaches can be susceptible to catastrophic forgetting whereas GPAs can easily incorporate new training data.

Our approach differs from existing work in several aspects. Our approach constructs a domain-dependent GPA automatically without any human intervention. Using canonical abstraction, we lift problem-specific characteristics like object names and object counts in a domain-independent fashion. Another key difference between other techniques is that our approach can easily incorporate solutions from new examples into the GPA without having to remember any of the earlier examples. This allows our learning to scale better and can naturally utilize *leapfrog learning* (Groshev et al., 2018; Karia and Srivastava, 2021, 2022) when presented with a large problem in the absence of training data. Finally, our approach comes with guarantees of completeness and hierarchical optimality given the training data presented. This means that if a solution exists, our approach will find it and it will be guaranteed to be hierarchically optimal or better (Thm. 3).

## 6 Conclusions and Future Work

We show that non-deterministic Generalized Policy Automata (GPAs) constructed using solutions of small example SSPs are able to significantly reduce the computational effort for finding solutions for larger related SSPs. Furthermore, for many benchmark problems, the search space pruned by GPAs does not prune away relevant transitions thus allowing our approach to compute policies of comparable cost in a fraction of the effort. Our approach comes with guarantees of hierarchical optimality and also comes with the guarantee of always finding a solution to the SSP.

There are several interesting research directions for future work. Description Logics are more expressive than canonical abstractions and have been demonstrated to be effective at synthesizing memoryless controllers for deterministic planning problems (Bonet et al., 2019). Our approach can easily utilize any relational abstraction and it would be interesting to evaluate the efficacy of description logics. Finally, our approach is applicable when solutions have a pattern. We believe that more intelligent training data generation methods could help improve performance in domains like Delicate Can. We plan to investigate these directions of research in future work.

## Acknowledgements

We would like to thank Deepak Kala Vasudevan for help with a prototype implementation of the source code. We would like to thank the Research Computing Group at Arizona State University for providing compute hours for our experiments. This work was supported in part by the NSF under grants IIS 1909370 and IIS 1942856.

## References

- J. S. Aguas, S. J. Celorrio, and A. Jonsson. Hierarchical finite state controllers for generalized planning. In *Proc. IJCAI*, 2016.
- A. Bai and S. J. Russell. Efficient reinforcement learning with hierarchies of machines by leveraging internal transitions. In *Proc. IJCAI*, 2017.
- D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-dynamic programming*. Athena Scientific, 1996. ISBN 1886529108.
- B. Bonet and H. Geffner. Labeled RTDP: improving the convergence of real-time dynamic programming. In *Proc. ICAPS*, 2003.
- B. Bonet, H. Palacios, and H. Geffner. Automatic derivation of memoryless policies and finite-state controllers using classical planners. In *Proc. ICAPS*, 2009.
- B. Bonet, G. Francès, and H. Geffner. Learning features and abstract actions for computing generalized plans. In *Proc. AAAI*, 2019.
- C. Boutilier, R. Reiter, and B. Price. Symbolic dynamic programming for first-order mdps. In *Proc. IJCAI*, 2001.
- A. Fern, S. W. Yoon, and R. Givan. Approximate policy iteration with a policy language bias: Solving relational markov decision processes. *J. Artif. Intell. Res.*, 25:75–118, 2006.
- G. Francès, B. Bonet, and H. Geffner. Learning general planning policies from small examples without supervision. In *Proc. AAAI*, 2021.
- S. Garg, A. Bajpai, and Mausam. Symbolic network: Generalized neural policies for relational mdps. In *Proc. ICML*, 2020.
- E. Groshev, M. Goldstein, A. Tamar, S. Srivastava, and P. Abbeel. Learning generalized reactive policies using deep neural networks. In *Proc. ICAPS*, 2018.
- E. A. Hansen and S. Zilberstein. Lao: A heuristic search algorithm that finds solutions with loops. *Artificial Intelligence*, 129(1-2):35–62, 2001.
- J. Hoffmann. FF: the fast-forward planning system. *AI Mag.*, 22(3):57–62, 2001.
- R. Karia and S. Srivastava. Learning generalized relational heuristic networks for model-agnostic planning. In *Proc. AAAI*, 2021.
- R. Karia and S. Srivastava. Relational abstractions for generalized reinforcement learning on symbolic problems. In *Proc. IJCAI*, 2022.
- R. Karia, R. K. Nayyar, and S. Srivastava. Learning generalized policy automata for relational stochastic shortest path problems. *arXiv*, abs/2204.04301, 2022.
- A. Kolobov, Mausam, and D. S. Weld. A theory of goal-oriented mdps with dead ends. In *Proc. UAI*, 2012.
- D. Long and M. Fox. The 3rd international planning competition: Results and analysis. *J. Artif. Intell. Res.*, 20:1–59, 2003.
- R. Parr and S. J. Russell. Reinforcement learning with hierarchies of machines. In *Proc. NeurIPS*, 1997.
- L. E. Pineda and S. Zilberstein. Planning under uncertainty using reduced models: Revisiting determinization. In *Proc. ICAPS*, 2014.
- L. E. Pineda and S. Zilberstein. Soft labeling in stochastic shortest path problems. In *Proc. AAMAS*, 2019.
- S. Sagiv, T. W. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *ACM Trans. Program. Lang. Syst.*, 24(3):217–298, 2002.

- S. Sanner. Relational dynamic influence diagram language (RDDL): Language description. [http://users.cecs.anu.edu.au/~ssanner/IPPC\\_2011/RDDL.pdf](http://users.cecs.anu.edu.au/~ssanner/IPPC_2011/RDDL.pdf), 2010.
- S. Sanner and C. Boutilier. Approximate linear programming for first-order mdps. In *Proc. UAI*, 2005.
- N. Shah, D. K. Vasudevan, K. Kumar, P. Kamojjhala, and S. Srivastava. Anytime integrated task and motion policies for stochastic environments. In *Proc. ICRA*, 2020.
- S. Srivastava, N. Immerman, and S. Zilberstein. Learning generalized plans using abstract counting. In *Proc. AAAI*, 2008.
- S. Srivastava, N. Immerman, and S. Zilberstein. A new representation and associated algorithms for generalized planning. *Artif. Intell.*, 175(2):615–647, 2011.
- S. Srivastava, N. Immerman, and S. Zilberstein. Applicability conditions for plans with loops: Computability results and algorithms. *Artif. Intell.*, 191-192:1–19, 2012.
- R. S. Sutton and A. G. Barto. *Reinforcement learning - an introduction*. MIT Press, 1998. ISBN 978-0-262-19398-6.
- A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel. Value iteration networks. In *Proc. NeurIPS*, 2016.
- S. Toyer, F. W. Trevizan, S. Thiébaux, and L. Xie. Action schema networks: Generalised policies with deep learning. In *Proc. AAAI*, 2018.
- F. W. Trevizan and M. M. Veloso. Short-sighted stochastic shortest path problems. In *Proc. ICAPS*, 2012.
- J. Wu and R. Givan. Discovering relational domain features for probabilistic planning. In *Proc. ICAPS*, 2007.
- S. W. Yoon, A. Fern, and R. Givan. Ff-replan: A baseline for probabilistic planning. In *Proc. ICAPS*, 2007.
- H. L. S. Younes, M. L. Littman, D. Weissman, and J. Asmuth. The first probabilistic track of the international planning competition. *J. Artif. Intell. Res.*, 24:851–887, 2005.

## A Extended Experiments and Results

**Training and Test Setup** Descriptions of the training problems used by us and their parameters can be found in Table. 1. Test problems and parameters along with complete information for the solution times, costs, and their standard deviations for 10 runs are available in tabular format in Tables 2, 3, 4, 5, and 6. Note that for the Keva domain, the standard deviations for the costs incurred are 0. This is accurate since the only source of stochasticity in Keva is a *human place* action that determines where the human places a plank which is one of two locations. As a result, Keva policies are deterministic in execution since the human always places a plank and all other actions are deterministic. It is interesting that despite this simplistic domain, the baselines are unable to reasonably converge within the timeout. We also present an extended version of Fig. 3 of the main paper that includes results for a larger suite a test problems for a better view of our overall results. These results are reported in Fig. 4. Note that the solution times on the left y-axis of these plots are shown in log scale. We omitted the problems with smaller IDs, mainly whose solutions times (in log scale) were not visible for both our as well as baseline approaches, for better visualization.

**Hyperparameters** We used  $\epsilon = 10^{-5}$  as the value for determining whether an algorithm has converged to an  $\epsilon$ -consistent policy. We set the total number of trials for all algorithms to  $\infty$ . As a result, each algorithm would only return once it has converged or if the time limit has been exceeded. For Soft-FLARES, we used  $t = 4$  which controls the horizon of the sub-tree that is checked for being  $\epsilon$ -consistent during the labeling procedure. Our distance metric is the *step* function which simply counts the depth until the horizon is exceeded. For the selective sampling procedure, we used the *logistic* sampler configured with  $\alpha = 0.1$  and  $\beta = 0.9$ .

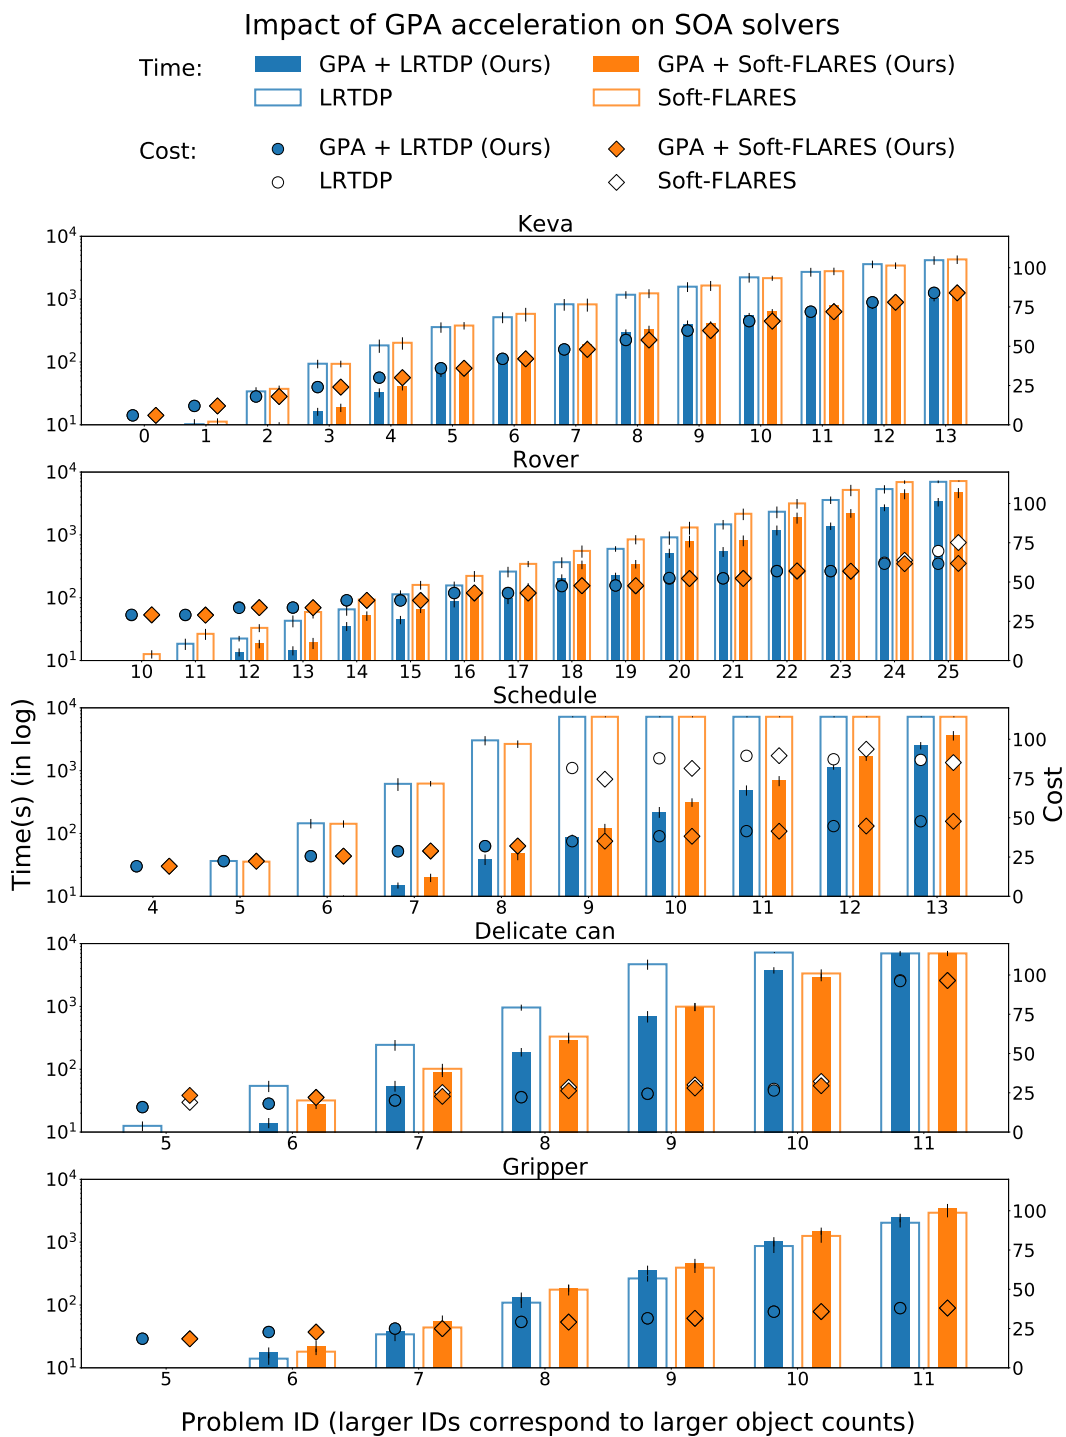


Figure 4: Impact of learned GPAs on solver performance of test problems (lower values better). Left y-axes and bars show solution times (in log scale), and right y-axes and points show cost incurred by the policy computed by our approach and baseline SOA solvers (LRTDP and Soft-FLARES). We use the same SSP solver as the corresponding baseline in our approach. Error bars for solution times indicate 1 standard deviation (SD) averaged across 10 runs.

Table 1: Our training setup for all domains. ID refers to the problem ID in the training set. The other columns refer to the parameters passed to the problem generator for generating the problem. Entries marked – indicate that there was no such problem ID in the training set.

ID	$\theta$				
	Keva( $P, h$ )	Rover( $r, w, s, o$ )	Schedule( $C, p$ )	Delicate Can( $c$ )	Gripper( $b$ )
0	(2, 1)	(1, 3, 1, 2)	(1, 2)	(2)	(1)
1	(4, 2)	(1, 4, 1, 2)	(1, 3)	(3)	(1)
2	(6, 3)	(1, 3, 2, 2)	(1, 4)	(4)	(2)
3	(8, 4)	(1, 4, 2, 2)	–	(5)	(2)
4	(10, 5)	(1, 3, 3, 2)	–	(6)	(3)
5	(12, 6)	(1, 4, 3, 2)	–	–	(3)
6	–	(1, 3, 4, 2)	–	–	(4)
7	–	(1, 4, 4, 2)	–	–	(4)
8	–	(1, 3, 5, 2)	–	–	(5)
9	–	(1, 4, 5, 2)	–	–	(5)

Table 2: Our test setup for the Keva( $P, h$ ) domain (lower values better). ID refers to the problem ID in the test set.  $\theta$  refers to the parameters passed to the problem generator for generating the problem. Times indicate the seconds required to find a policy. Similarly, costs are reported as average costs obtained by executing the computed policy for 100 trials. We ran our experiments using a different random seed for 10 different runs and report average metrics up to one standard deviation. Better metrics are at least 5% better and are indicated using bold font.

ID	$\theta$	Time( $x \equiv$ LRTDP)		Time( $x \equiv$ Soft-FLARES)		Cost( $x \equiv$ LRTDP)		Cost( $x \equiv$ Soft-FLARES)	
		$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$
0	(29, 1)	2.45 ±0.32	<b>1.44</b> ±0.17	2.65 ±0.37	<b>1.47</b> ±0.27	6.00 ±0.00	6.00 ±0.00	6.00 ±0.00	6.00 ±0.00
1	(29, 2)	10.23 ±2.05	<b>4.15</b> ±0.68	11.22 ±1.45	<b>4.63</b> ±0.58	12.00 ±0.00	12.00 ±0.00	12.00 ±0.00	12.00 ±0.00
2	(29, 3)	34.05 ±5.53	<b>8.14</b> ±0.92	37.33 ±4.71	<b>9.55</b> ±1.45	18.00 ±0.00	18.00 ±0.00	18.00 ±0.00	18.00 ±0.00
3	(29, 4)	93.78 ±14.53	<b>16.26</b> ±2.34	93.39 ±11.30	<b>18.80</b> ±2.84	24.00 ±0.00	24.00 ±0.00	24.00 ±0.00	24.00 ±0.00
4	(29, 5)	183.57 ±43.80	<b>32.61</b> ±5.40	201.99 ±45.05	<b>41.19</b> ±6.05	30.00 ±0.00	30.00 ±0.00	30.00 ±0.00	30.00 ±0.00
5	(29, 6)	358.98 ±67.32	<b>68.73</b> ±11.04	379.84 ±51.43	<b>74.47</b> ±9.10	36.00 ±0.00	36.00 ±0.00	36.00 ±0.00	36.00 ±0.00
6	(29, 7)	515.85 ±101.56	<b>115.28</b> ±11.90	583.52 ±143.15	<b>120.30</b> ±17.58	42.00 ±0.00	42.00 ±0.00	42.00 ±0.00	42.00 ±0.00
7	(29, 8)	829.19 ±174.46	<b>174.88</b> ±17.36	825.84 ±192.54	<b>191.44</b> ±27.16	48.00 ±0.00	48.00 ±0.00	48.00 ±0.00	48.00 ±0.00
8	(29, 9)	1174.79 ±160.45	<b>298.61</b> ±30.82	1236.91 ±200.02	<b>331.95</b> ±45.28	54.00 ±0.00	54.00 ±0.00	54.00 ±0.00	54.00 ±0.00
9	(29, 10)	1578.38 ±279.31	<b>394.71</b> ±62.16	1647.76 ±297.42	<b>406.45</b> ±37.23	60.00 ±0.00	60.00 ±0.00	60.00 ±0.00	60.00 ±0.00
10	(29, 11)	2223.43 ±390.05	<b>544.26</b> ±58.84	2158.91 ±199.54	<b>639.60</b> ±54.35	66.00 ±0.00	66.00 ±0.00	66.00 ±0.00	66.00 ±0.00
11	(29, 12)	2713.78 ±435.39	<b>665.29</b> ±76.02	2787.92 ±378.85	<b>782.82</b> ±74.38	72.00 ±0.00	72.00 ±0.00	72.00 ±0.00	72.00 ±0.00
12	(29, 13)	3606.12 ±494.69	<b>815.11</b> ±117.62	3427.93 ±409.34	<b>958.17</b> ±81.59	78.00 ±0.00	78.00 ±0.00	78.00 ±0.00	78.00 ±0.00
13	(29, 14)	4171.27 ±645.70	<b>1042.44</b> ±119.63	4291.64 ±663.83	<b>1128.84</b> ±183.24	84.00 ±0.00	84.00 ±0.00	84.00 ±0.00	84.00 ±0.00

Table 3: Our test setup for the Rover( $r, w, s, o$ ) domain (lower values better). ID refers to the problem ID in the test set.  $\theta$  refers to the parameters passed to the problem generator for generating the problem. Times indicate the seconds required to find a policy. Similarly, costs are reported as average costs obtained by executing the computed policy for 100 trials. We ran our experiments using a different random seed for 10 different runs and report average metrics up to one standard deviation. Better metrics are at least 5% better and are indicated using bold font.

ID	$\theta$	Time( $x \equiv$ LRTDP)		Time( $x \equiv$ Soft-FLARES)		Cost( $x \equiv$ LRTDP)		Cost( $x \equiv$ Soft-FLARES)	
		$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$
0	(1, 3, 1, 2)	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	<b>0.01</b> $\pm$ 0.00	6.62 $\pm$ 0.08	6.68 $\pm$ 0.08	6.69 $\pm$ 0.12	6.68 $\pm$ 0.09
1	(1, 4, 1, 2)	0.02 $\pm$ 0.01	0.02 $\pm$ 0.00	0.02 $\pm$ 0.01	<b>0.01</b> $\pm$ 0.00	6.64 $\pm$ 0.08	6.70 $\pm$ 0.09	6.62 $\pm$ 0.12	6.70 $\pm$ 0.12
2	(1, 3, 2, 2)	0.10 $\pm$ 0.03	<b>0.03</b> $\pm$ 0.01	0.09 $\pm$ 0.02	<b>0.03</b> $\pm$ 0.00	10.33 $\pm$ 0.14	10.29 $\pm$ 0.08	10.38 $\pm$ 0.15	10.34 $\pm$ 0.10
3	(1, 4, 2, 2)	0.18 $\pm$ 0.04	<b>0.03</b> $\pm$ 0.01	0.22 $\pm$ 0.02	<b>0.04</b> $\pm$ 0.01	10.34 $\pm$ 0.11	10.30 $\pm$ 0.12	10.27 $\pm$ 0.12	10.26 $\pm$ 0.16
4	(1, 3, 3, 2)	0.43 $\pm$ 0.06	<b>0.18</b> $\pm$ 0.03	0.46 $\pm$ 0.07	<b>0.20</b> $\pm$ 0.04	15.03 $\pm$ 0.19	15.01 $\pm$ 0.16	14.94 $\pm$ 0.23	15.09 $\pm$ 0.11
5	(1, 4, 3, 2)	0.80 $\pm$ 0.14	<b>0.10</b> $\pm$ 0.02	0.96 $\pm$ 0.24	<b>0.09</b> $\pm$ 0.01	15.00 $\pm$ 0.23	15.01 $\pm$ 0.23	14.95 $\pm$ 0.17	14.95 $\pm$ 0.19
6	(1, 3, 4, 2)	1.08 $\pm$ 0.12	<b>0.54</b> $\pm$ 0.10	1.66 $\pm$ 0.39	<b>0.68</b> $\pm$ 0.11	19.76 $\pm$ 0.21	19.65 $\pm$ 0.31	19.68 $\pm$ 0.15	19.62 $\pm$ 0.17
7	(1, 4, 4, 2)	2.33 $\pm$ 0.43	<b>0.54</b> $\pm$ 0.08	3.22 $\pm$ 0.69	<b>0.70</b> $\pm$ 0.14	19.63 $\pm$ 0.13	19.63 $\pm$ 0.18	19.70 $\pm$ 0.16	19.73 $\pm$ 0.20
8	(1, 3, 5, 2)	3.54 $\pm$ 0.59	<b>1.61</b> $\pm$ 0.23	4.25 $\pm$ 0.84	<b>2.42</b> $\pm$ 0.51	24.36 $\pm$ 0.17	24.35 $\pm$ 0.30	24.34 $\pm$ 0.22	24.46 $\pm$ 0.09
9	(1, 4, 5, 2)	7.78 $\pm$ 1.62	<b>1.82</b> $\pm$ 0.28	9.57 $\pm$ 1.69	<b>2.26</b> $\pm$ 0.42	24.41 $\pm$ 0.27	24.38 $\pm$ 0.16	24.32 $\pm$ 0.34	24.23 $\pm$ 0.32
10	(1, 3, 6, 2)	8.77 $\pm$ 1.08	<b>4.91</b> $\pm$ 1.04	12.67 $\pm$ 1.91	<b>6.40</b> $\pm$ 0.93	28.95 $\pm$ 0.26	29.13 $\pm$ 0.26	28.98 $\pm$ 0.25	29.14 $\pm$ 0.31
11	(1, 4, 6, 2)	18.50 $\pm$ 3.76	<b>4.31</b> $\pm$ 0.77	26.62 $\pm$ 5.28	<b>5.28</b> $\pm$ 0.56	28.96 $\pm$ 0.29	29.07 $\pm$ 0.18	29.09 $\pm$ 0.18	28.86 $\pm$ 0.25
12	(1, 3, 7, 2)	22.42 $\pm$ 2.26	<b>13.60</b> $\pm$ 1.98	33.11 $\pm$ 4.79	<b>18.49</b> $\pm$ 2.85	33.80 $\pm$ 0.29	33.54 $\pm$ 0.20	33.74 $\pm$ 0.26	33.71 $\pm$ 0.24
13	(1, 4, 7, 2)	42.96 $\pm$ 9.15	<b>14.47</b> $\pm$ 2.40	59.22 $\pm$ 12.54	<b>19.09</b> $\pm$ 3.84	33.70 $\pm$ 0.34	33.75 $\pm$ 0.22	33.63 $\pm$ 0.21	33.67 $\pm$ 0.28
14	(1, 3, 8, 2)	65.32 $\pm$ 13.63	<b>35.25</b> $\pm$ 5.75	93.44 $\pm$ 13.87	<b>51.72</b> $\pm$ 8.90	38.30 $\pm$ 0.26	38.39 $\pm$ 0.24	38.28 $\pm$ 0.27	38.36 $\pm$ 0.35
15	(1, 4, 8, 2)	113.36 $\pm$ 17.49	<b>44.60</b> $\pm$ 6.77	159.83 $\pm$ 25.30	<b>64.87</b> $\pm$ 7.88	38.43 $\pm$ 0.17	38.18 $\pm$ 0.22	38.33 $\pm$ 0.29	38.24 $\pm$ 0.36
16	(1, 3, 9, 2)	156.88 $\pm$ 23.55	<b>86.26</b> $\pm$ 16.11	222.42 $\pm$ 43.47	<b>122.44</b> $\pm$ 21.53	43.02 $\pm$ 0.40	42.98 $\pm$ 0.45	42.87 $\pm$ 0.26	43.11 $\pm$ 0.29
17	(1, 4, 9, 2)	260.78 $\pm$ 49.61	<b>95.92</b> $\pm$ 16.68	345.68 $\pm$ 38.97	<b>142.71</b> $\pm$ 27.75	43.05 $\pm$ 0.19	42.94 $\pm$ 0.31	42.98 $\pm$ 0.36	43.00 $\pm$ 0.24
18	(1, 3, 10, 2)	367.77 $\pm$ 71.46	<b>199.80</b> $\pm$ 35.50	555.65 $\pm$ 121.90	<b>337.10</b> $\pm$ 52.84	47.65 $\pm$ 0.34	47.36 $\pm$ 0.16	47.62 $\pm$ 0.30	47.67 $\pm$ 0.42
19	(1, 4, 10, 2)	599.74 $\pm$ 60.23	<b>223.38</b> $\pm$ 26.79	848.57 $\pm$ 141.78	<b>345.02</b> $\pm$ 53.64	47.85 $\pm$ 0.39	47.78 $\pm$ 0.27	47.49 $\pm$ 0.28	47.70 $\pm$ 0.19
20	(1, 3, 11, 2)	914.39 $\pm$ 217.86	<b>515.34</b> $\pm$ 85.61	1312.81 $\pm$ 302.78	<b>800.94</b> $\pm$ 169.25	52.17 $\pm$ 0.22	52.58 $\pm$ 0.35	52.29 $\pm$ 0.32	52.21 $\pm$ 0.29
21	(1, 4, 11, 2)	1472.73 $\pm$ 254.76	<b>543.15</b> $\pm$ 97.44	2168.73 $\pm$ 450.03	<b>819.44</b> $\pm$ 159.47	52.36 $\pm$ 0.27	52.35 $\pm$ 0.41	52.37 $\pm$ 0.46	52.26 $\pm$ 0.30
22	(1, 3, 12, 2)	2336.52 $\pm$ 492.69	<b>1195.16</b> $\pm$ 213.55	3171.35 $\pm$ 554.37	<b>1885.02</b> $\pm$ 375.15	57.00 $\pm$ 0.35	56.94 $\pm$ 0.26	56.85 $\pm$ 0.36	57.12 $\pm$ 0.62
23	(1, 4, 12, 2)	3593.42 $\pm$ 487.19	<b>1385.19</b> $\pm$ 187.00	5196.66 $\pm$ 1063.94	<b>2224.24</b> $\pm$ 351.60	56.93 $\pm$ 0.24	56.94 $\pm$ 0.24	56.80 $\pm$ 0.27	56.94 $\pm$ 0.35
24	(1, 3, 13, 2)	5366.69 $\pm$ 807.68	<b>2721.86</b> $\pm$ 337.95	6933.16 $\pm$ 432.32	<b>4505.52</b> $\pm$ 811.02	62.26 $\pm$ 1.69	61.63 $\pm$ 0.39	63.83 $\pm$ 2.43	61.67 $\pm$ 0.29
25	(1, 4, 13, 2)	6997.37 $\pm$ 338.99	<b>3349.16</b> $\pm$ 517.48	7200.00 $\pm$ 0.00	<b>4710.61</b> $\pm$ 855.90	69.72 $\pm$ 11.30	<b>61.60</b> $\pm$ 0.46	75.14 $\pm$ 13.01	<b>61.83</b> $\pm$ 0.37



Table 4: Our test setup for the Schedule( $C, p$ ) domain (lower values better). ID refers to the problem ID in the test set.  $\theta$  refers to the parameters passed to the problem generator for generating the problem. Times indicate the seconds required to find a policy. Similarly, costs are reported as average costs obtained by executing the computed policy for 100 trials. We ran our experiments using a different random seed for 10 different runs and report average metrics up to one standard deviation. Better metrics are at least 5% better and are indicated using bold font.

ID	$\theta$	Time( $x \equiv$ LRTDP)		Time( $x \equiv$ Soft-FLARES)		Cost( $x \equiv$ LRTDP)		Cost( $x \equiv$ Soft-FLARES)	
		$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$
0	(1, 2)	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	6.41 $\pm$ 0.13	6.32 $\pm$ 0.13	6.35 $\pm$ 0.08	6.45 $\pm$ 0.15
1	(1, 3)	0.08 $\pm$ 0.02	<b>0.06</b> $\pm$ 0.02	0.07 $\pm$ 0.02	0.07 $\pm$ 0.02	9.53 $\pm$ 0.12	9.50 $\pm$ 0.09	9.60 $\pm$ 0.13	9.52 $\pm$ 0.12
2	(1, 4)	0.32 $\pm$ 0.05	<b>0.16</b> $\pm$ 0.03	0.33 $\pm$ 0.05	<b>0.19</b> $\pm$ 0.03	12.72 $\pm$ 0.14	12.73 $\pm$ 0.14	12.81 $\pm$ 0.15	12.77 $\pm$ 0.15
3	(1, 5)	1.58 $\pm$ 0.25	<b>0.41</b> $\pm$ 0.06	1.60 $\pm$ 0.34	<b>0.46</b> $\pm$ 0.09	16.00 $\pm$ 0.14	15.97 $\pm$ 0.17	15.89 $\pm$ 0.11	15.95 $\pm$ 0.19
4	(1, 6)	6.45 $\pm$ 0.77	<b>1.02</b> $\pm$ 0.22	7.36 $\pm$ 1.28	<b>1.21</b> $\pm$ 0.20	19.06 $\pm$ 0.16	19.16 $\pm$ 0.16	19.17 $\pm$ 0.24	19.14 $\pm$ 0.14
5	(1, 7)	36.46 $\pm$ 7.19	<b>2.46</b> $\pm$ 0.56	35.61 $\pm$ 6.35	<b>3.10</b> $\pm$ 0.59	22.37 $\pm$ 0.19	22.45 $\pm$ 0.19	22.45 $\pm$ 0.25	22.28 $\pm$ 0.24
6	(1, 8)	145.33 $\pm$ 24.97	<b>6.58</b> $\pm$ 1.25	142.70 $\pm$ 18.86	<b>8.42</b> $\pm$ 1.93	25.57 $\pm$ 0.18	25.56 $\pm$ 0.09	25.52 $\pm$ 0.23	25.59 $\pm$ 0.26
7	(1, 9)	616.36 $\pm$ 140.89	<b>14.92</b> $\pm$ 1.65	622.93 $\pm$ 61.96	<b>19.85</b> $\pm$ 3.00	28.78 $\pm$ 0.18	28.61 $\pm$ 0.16	28.73 $\pm$ 0.25	28.88 $\pm$ 0.18
8	(1, 10)	3036.01 $\pm$ 507.41	<b>38.89</b> $\pm$ 7.41	2662.96 $\pm$ 361.97	<b>48.08</b> $\pm$ 10.54	31.95 $\pm$ 0.08	31.95 $\pm$ 0.23	31.96 $\pm$ 0.20	32.01 $\pm$ 0.21
9	(1, 11)	7200.00 $\pm$ 0.00	<b>85.99</b> $\pm$ 9.94	7200.00 $\pm$ 0.00	<b>122.28</b> $\pm$ 19.92	81.68 $\pm$ 16.65	<b>35.06</b> $\pm$ 0.29	74.56 $\pm$ 19.98	<b>35.10</b> $\pm$ 0.28
10	(1, 12)	7200.00 $\pm$ 0.00	<b>220.64</b> $\pm$ 43.34	7200.00 $\pm$ 0.00	<b>313.88</b> $\pm$ 48.81	87.85 $\pm$ 11.06	<b>38.26</b> $\pm$ 0.23	81.42 $\pm$ 16.54	<b>38.20</b> $\pm$ 0.25
11	(1, 13)	7200.00 $\pm$ 0.00	<b>490.90</b> $\pm$ 90.83	7200.00 $\pm$ 0.00	<b>692.75</b> $\pm$ 126.60	89.45 $\pm$ 12.88	<b>41.50</b> $\pm$ 0.40	89.54 $\pm$ 11.68	<b>41.43</b> $\pm$ 0.22
12	(1, 14)	7200.00 $\pm$ 0.00	<b>1153.32</b> $\pm$ 117.37	7200.00 $\pm$ 0.00	<b>1710.56</b> $\pm$ 278.99	87.26 $\pm$ 12.32	<b>44.69</b> $\pm$ 0.13	93.66 $\pm$ 10.10	<b>44.74</b> $\pm$ 0.27
13	(1, 15)	7200.00 $\pm$ 0.00	<b>2514.16</b> $\pm$ 337.59	7200.00 $\pm$ 0.00	<b>3639.92</b> $\pm$ 627.08	86.84 $\pm$ 11.41	<b>47.77</b> $\pm$ 0.17	85.14 $\pm$ 13.60	<b>47.74</b> $\pm$ 0.34

Table 5: Our test setup for the Delicate Can( $c$ ) domain (lower values better). ID refers to the problem ID in the test set.  $\theta$  refers to the parameters passed to the problem generator for generating the problem. Times indicate the seconds required to find a policy. Similarly, costs are reported as average costs obtained by executing the computed policy for 100 trials. We ran our experiments using a different random seed for 10 different runs and report average metrics up to one standard deviation. Better metrics are at least 5% better and are indicated using bold font.

ID	$\theta$	Time( $x \equiv$ LRTDP)		Time( $x \equiv$ Soft-FLARES)		Cost( $x \equiv$ LRTDP)		Cost( $x \equiv$ Soft-FLARES)	
		$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$
0	(2)	0.01 $\pm$ 0.01	0.01 $\pm$ 0.00	<b>0.00</b> $\pm$ 0.00	0.01 $\pm$ 0.01	5.40 $\pm$ 0.07	5.40 $\pm$ 0.08	5.45 $\pm$ 0.07	5.40 $\pm$ 0.10
1	(3)	0.03 $\pm$ 0.01	<b>0.02</b> $\pm$ 0.01	0.03 $\pm$ 0.01	<b>0.02</b> $\pm$ 0.01	7.42 $\pm$ 0.12	7.50 $\pm$ 0.15	7.53 $\pm$ 0.09	7.48 $\pm$ 0.06
2	(4)	0.11 $\pm$ 0.03	<b>0.07</b> $\pm$ 0.02	0.12 $\pm$ 0.03	<b>0.09</b> $\pm$ 0.02	9.58 $\pm$ 0.11	9.53 $\pm$ 0.16	9.56 $\pm$ 0.14	9.62 $\pm$ 0.09
3	(5)	0.49 $\pm$ 0.08	<b>0.22</b> $\pm$ 0.05	0.51 $\pm$ 0.08	<b>0.40</b> $\pm$ 0.06	11.71 $\pm$ 0.20	11.72 $\pm$ 0.12	<b>12.48</b> $\pm$ 0.69	14.22 $\pm$ 2.49
4	(6)	2.56 $\pm$ 0.55	<b>0.80</b> $\pm$ 0.13	2.14 $\pm$ 0.26	<b>1.79</b> $\pm$ 0.23	13.82 $\pm$ 0.15	13.75 $\pm$ 0.07	<b>15.02</b> $\pm$ 1.09	25.16 $\pm$ 4.14
5	(7)	12.55 $\pm$ 2.20	<b>3.25</b> $\pm$ 0.56	8.08 $\pm$ 0.92	<b>6.81</b> $\pm$ 1.41	15.94 $\pm$ 0.09	15.86 $\pm$ 0.08	<b>18.87</b> $\pm$ 1.44	23.32 $\pm$ 4.09
6	(8)	54.31 $\pm$ 11.04	<b>14.12</b> $\pm$ 2.60	31.86 $\pm$ 5.15	<b>27.65</b> $\pm$ 4.41	18.08 $\pm$ 0.13	18.07 $\pm$ 0.14	22.16 $\pm$ 1.68	22.01 $\pm$ 1.52
7	(9)	244.32 $\pm$ 47.46	<b>54.60</b> $\pm$ 10.68	101.63 $\pm$ 20.29	<b>90.57</b> $\pm$ 15.57	20.04 $\pm$ 0.13	20.15 $\pm$ 0.14	25.18 $\pm$ 2.65	<b>22.70</b> $\pm$ 1.22
8	(10)	960.96 $\pm$ 105.78	<b>188.50</b> $\pm$ 29.21	331.06 $\pm$ 49.98	<b>293.54</b> $\pm$ 36.06	22.28 $\pm$ 0.17	22.17 $\pm$ 0.13	28.54 $\pm$ 2.09	<b>26.29</b> $\pm$ 1.73
9	(11)	4691.23 $\pm$ 859.08	<b>698.09</b> $\pm$ 143.63	990.65 $\pm$ 144.48	983.51 $\pm$ 144.08	24.30 $\pm$ 0.12	24.41 $\pm$ 0.17	30.05 $\pm$ 1.37	<b>28.02</b> $\pm$ 1.40
10	(12)	7200.00 $\pm$ 0.00	<b>3769.57</b> $\pm$ 439.21	3347.02 $\pm$ 540.66	<b>2911.30</b> $\pm$ 397.41	27.46 $\pm$ 0.83	26.46 $\pm$ 0.14	32.23 $\pm$ 2.35	<b>29.49</b> $\pm$ 1.47
11	(12)	6980.00 $\pm$ 660.00	7183.58 $\pm$ 49.27	6980.48 $\pm$ 658.56	7132.65 $\pm$ 202.04	96.62 $\pm$ 10.13	96.12 $\pm$ 11.65	96.50 $\pm$ 10.49	96.48 $\pm$ 10.56

Table 6: Our test setup for the Gripper( $b$ ) domain (lower values better). ID refers to the problem ID in the test set.  $\theta$  refers to the parameters passed to the problem generator for generating the problem. Times indicate the seconds required to find a policy. Similarly, costs are reported as average costs obtained by executing the computed policy for 100 trials. We ran our experiments using a different random seed for 10 different runs and report average metrics up to one standard deviation. Better metrics are at least 5% better and are indicated using bold font.

ID	$\theta$	Time( $x \equiv$ LRTDP)		Time( $x \equiv$ Soft-FLARES)		Cost( $x \equiv$ LRTDP)		Cost( $x \equiv$ Soft-FLARES)	
		$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$	$x$	Ours + $x$
0	(1)	0.01 $\pm$ 0.01	<b>0.00</b> $\pm$ 0.01	<b>0.00</b> $\pm$ 0.00	0.01 $\pm$ 0.01	3.25 $\pm$ 0.06	3.25 $\pm$ 0.05	3.23 $\pm$ 0.06	3.25 $\pm$ 0.04
1	(2)	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	0.02 $\pm$ 0.01	5.48 $\pm$ 0.08	5.53 $\pm$ 0.08	5.51 $\pm$ 0.10	5.55 $\pm$ 0.07
2	(3)	<b>0.10</b> $\pm$ 0.02	0.12 $\pm$ 0.03	0.12 $\pm$ 0.03	0.12 $\pm$ 0.03	9.71 $\pm$ 0.07	9.76 $\pm$ 0.10	9.76 $\pm$ 0.07	9.75 $\pm$ 0.09
3	(4)	0.33 $\pm$ 0.05	0.34 $\pm$ 0.05	<b>0.30</b> $\pm$ 0.06	0.38 $\pm$ 0.07	11.99 $\pm$ 0.10	12.03 $\pm$ 0.16	12.00 $\pm$ 0.08	12.03 $\pm$ 0.10
4	(5)	<b>1.36</b> $\pm$ 0.23	1.62 $\pm$ 0.28	<b>1.82</b> $\pm$ 0.35	2.21 $\pm$ 0.47	16.27 $\pm$ 0.09	16.23 $\pm$ 0.12	16.22 $\pm$ 0.06	16.26 $\pm$ 0.15
5	(6)	<b>3.38</b> $\pm$ 0.39	5.17 $\pm$ 1.65	<b>4.64</b> $\pm$ 0.80	5.16 $\pm$ 0.64	18.51 $\pm$ 0.21	18.54 $\pm$ 0.14	18.53 $\pm$ 0.12	18.44 $\pm$ 0.07
6	(7)	<b>14.00</b> $\pm$ 2.83	17.67 $\pm$ 3.43	<b>18.07</b> $\pm$ 2.13	22.28 $\pm$ 4.48	22.82 $\pm$ 0.08	22.78 $\pm$ 0.10	22.80 $\pm$ 0.19	22.73 $\pm$ 0.15
7	(8)	<b>34.11</b> $\pm$ 7.52	38.35 $\pm$ 4.68	<b>43.88</b> $\pm$ 8.90	55.78 $\pm$ 12.25	24.96 $\pm$ 0.12	24.99 $\pm$ 0.13	24.96 $\pm$ 0.14	24.95 $\pm$ 0.14
8	(9)	<b>108.94</b> $\pm$ 19.36	133.06 $\pm$ 24.75	<b>175.48</b> $\pm$ 33.68	186.91 $\pm$ 24.64	29.17 $\pm$ 0.10	29.26 $\pm$ 0.20	29.23 $\pm$ 0.14	29.08 $\pm$ 0.17
9	(10)	<b>264.59</b> $\pm$ 28.50	362.16 $\pm$ 60.65	<b>392.91</b> $\pm$ 69.13	463.30 $\pm$ 78.81	31.48 $\pm$ 0.20	31.54 $\pm$ 0.15	31.47 $\pm$ 0.14	31.54 $\pm$ 0.17
10	(11)	<b>867.91</b> $\pm$ 195.44	1028.70 $\pm$ 172.62	<b>1256.52</b> $\pm$ 277.94	1509.25 $\pm$ 195.14	35.78 $\pm$ 0.09	35.73 $\pm$ 0.15	35.70 $\pm$ 0.15	35.84 $\pm$ 0.16
11	(12)	<b>2037.65</b> $\pm$ 326.63	2454.41 $\pm$ 377.38	<b>2941.24</b> $\pm$ 457.95	3461.67 $\pm$ 595.74	38.07 $\pm$ 0.13	37.98 $\pm$ 0.15	38.03 $\pm$ 0.18	38.10 $\pm$ 0.12