Learning Generalized Policy Automata for Relational Stochastic Shortest Path Problems

Anonymous Author(s)
Affiliation
Address
email

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. Currently, existing state-of-the-art planners and heuristics often fail to exploit knowledge learned from solving other instances. 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 finding solutions to goal-oriented MDPs using image-based state representations (Tamar et al., 2016; Pong et al., 2018; Levy et al., 2019). However, these methods require significant human-engineering effort in finding transformations like grayscale conversion, etc., to yield representations that facilitate learning. A large number of practical problems however, are more intuitively expressed using relational representations and have been widely studied in the literature. 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 lends itself to a relational description. Converting these existing problems to suitable image-based representations would be counter-productive and difficult. Furthermore, they either require large amounts of training data and/or are unable to provide guarantees of completeness and/or convergence.

Many real-world problems can be readily expressed as SSPs and polynomial-time algorithms are available for solving them. A common theme among different SSP solvers is to 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. Despite such optimizations, a major hurdle that SSP solvers face is the "curse-of-dimensionality" since the state spaces grow exponentially as the total number of objects increases. The pruning strategies employed by these SSP solvers fail because they are "stateless" in the sense that they do not store, and consequently fail to re-use, knowledge that could have been exploited by solving similar problems.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

For rover problems, the total number of possible states grows exponentially as a factor of the total number of locations and rocks. Existing SSP solvers would have difficulty scaling to problems with many locations and/or rocks. A better way to tackle this problem of state space explosion would be to compute a simple *generalized policy*: move the rover to the closest available location with a rock, try loading the rock until it succeeds, navigate back to the base, unload it, and re-iterate this process until all the rocks are at the base.

Such a generalized policy can be used to solve many different SSP problems with large numbers of objects that share similar objectives. Furthermore, such policies appear to be easily computable using solutions of problems with few numbers of objects. Recent work has demonstrated that it is possible to learn such generalized policies. Toyer et al. (2018); Groshev et al. (2018b); Garg et al. (2020) use deep learning to learn generalized reactive policies. Bonet et al. (2009) automatically derive controllers that represent generalized reactive policies. A key limitation of these approaches is the lack of any theoretical guarantees of finding a solution or its optimality even if one is guaranteed to exist. 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 *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*. Our approach uses GPAs to prune out large sets of actions in different, related but larger SSPs. We prove that our approach is complete and guarantees hierarchical optimality. Empirical analysis on a range of well-known benchmark domains used in various planning competitions and robotics shows that our approach few-shot learns GPAs and convincingly outperforms existing state-of-the-art (SOA) 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 use 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 $\mathcal{P}$ is a set of predicates and $\mathcal{A}$ is a set of parameterized actions. Object *types*, such as those used in PPDDL, can be equivalently represented using unary predicates. A relational SSP problem for a domain $\mathcal{D}$ is then defined as a tuple $P = \langle O, S, A, s_0, g, T, C \rangle$ where O is a set of objects. 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 state s and ending up in a state s' where $a \in A$, $s, s' \in S$, 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 $a(o_1, \ldots, o_n)$, where $o_1, \ldots, o_n \in O$ are the action *parameters*. We omit the parameters when it is clear from context. $s_0 \in S$ is a *known* initial state and g is the goal condition expressed as a first-order logic formula. A goal state s_g is a state s.t. $s_g \models g$. $c(s_g, a, s_g) = 0$ and $t(s_g, a, s_g) = 1$ for all such goal states for any action a. Additionally, termination (reaching a state s.t. $s \models g$) in an SSP is inevitable making the length of the horizon unknown but finite (Bertsekas and Tsitsiklis, 1996).

Running example: The planetary rover example discussed earlier can be expressed using a domain that consists of parameterized predicates *connected*(l_x, l_y), *in-rover*(r_x), *rock-at*(r_x, l_x), and 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, allowing an easy way to express many different problems with different object names and counts. Actions dynamics are described using closed-form probability distributions (e.g. loading a rock could be modeled to fail with a probability of 0.2) 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 $O = \{l_1, l_2, \textit{base}, r_1, r_2\}$. 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 $s_{eg} = \{\textit{location}(l_1), \textit{location}(l_2), \textit{location}(\textit{base}), \textit{rock}(r_1), \textit{rock}(r_2), \textit{in-rover}(r_2), \textit{rock-at}(r_1, l_2)\}$. The goal of delivering all the rocks to the base can be expressed as $\forall r_x\ \textit{rock}(r_x) \wedge \textit{rock-at}(r_x, \textit{base})$. Actions are assumed to have a unit-cost when executed in a non-goal state and 0 otherwise.

A solution to an SSP is a deterministic policy $\pi : S \rightarrow A$ which is a mapping from states to actions. A *complete proper policy* is one for which the goal is guaranteed to be reachable 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. 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 from s_0 that is defined for every reachable state from s_0. To use such a formulation, we focus on a broader class of relaxed SSPs called Generalized SSPs (Kolobov et al., 2012) that allow the presence of dead-end states (whose costs are infinite) and only require the existence of at least one *partial* proper policy from s_0. Henceforth, we use the term SSPs to refer to Generalized SSPs and focus only on *partial* proper policies.

The value of a state s when using a policy π is the expected cost when starting in s and following π thereafter (Sutton and Barto, 1998):

$$v_\pi(s) = \sum_{s' \in S} t(s, \pi(s), s')[c(s, \pi(s), s') + v_\pi(s')]$$

This equation is the *Bellman* equation for the value of a state under the policy π. The *optimal* policy π_* is one s.t. $v_{\pi_*}(s) \leq v_\pi(s)$ for any policy π. The *Bellman optimality* equation (or the optimal state value function v_*) can be similarly expressed as (Sutton and Barto, 1998):

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

Once v_* has been computed, the optimal policy can be extracted using $\pi_*(s) = \arg\min_a v_*(s)$. Thus, given an SSP, the objective is to compute a policy that reaches the goal state from the initial state while minimizing the expected cost of doing so. SSP solvers often compute partial proper policies $\pi_*(s_0)$ by iteratively improving state-value estimates using the Bellman equations. Naturally, for any improper policy π, $v_\pi(s_0) = \infty$. SSP solvers have, under certain conditions, been proved to converge to a policy that is ϵ-consistent with the optimal policy (Hansen and Zilberstein, 2001; Bonet and Geffner, 2003).

A policy π for an SSP P can be represented as a concrete directed hypergraph $G_\pi = \langle V, E \rangle$ where $V = S$ is the vertex set comprising states of P and $E \subseteq V \times \mathcal{P}(V) \times A$ is the set of directed hyperedges s.t. each hyperedge $e \in E$ is a tuple $(e_{\textit{start}}, e_{\textit{results}}, e_{\textit{action}})$ representing the start vertex, a set of result vertices, and an action label. G_π can easily be constructed using the transition function T of P by adding vertices and hyperedges corresponding to each transition from T starting from $\pi(s_0)$. It is easy to see that such a hypergraph representation of π is a non-deterministic finite-state automaton whose states are bounded by the total number of vertices in G_π.

Let f be a feature and F be a set 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. We define *state abstraction* as a function $\alpha : F, S \rightarrow \overline{S}$ that transforms the concrete state space S for an SSP into an abstract state space $\overline{S}$ using the feature set F. Similarly, *action abstraction* $\beta : F, S, A \rightarrow \overline{A}$ transforms the action space to an abstract action space using the feature set F. In this paper, we utilize canonical abstraction (Sagiv et al., 2002) to compute such feature-based representations of concrete states and actions. We define feature sets in more detail in Sec. 3.2.

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 define GPAs and provide a brief description of canonical abstractions in Sec. 3.1 and Sec. 3.2 respectively, and describe our process to learn a GPA in Sec. 3.3. We then describe our method (Alg. 1) for solving SSPs in Sec. 3.4.

3.1 Generalized Policy Automata

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

Definition 3.1 (Abstract states and abstract actions)**.** Given a concrete directed hypergraph $G_\pi = \langle V, E \rangle$ that represents a policy π for an SSP P, abstraction functions α and β, and feature sets F_α and F_β, a set of abstract states $\overline{S}$ is defined as $\{\alpha(F_\alpha, s) | s \in V\}$. Similarly, a set of abstract actions $\overline{A}$ is defined as $\{\beta(F_\beta, s, a) | s \in V, e \in E, s \in e_{start}, a \in e_{action}\}$.

Definition 3.2 (Generalized Policy Automaton)**.** Given sets of abstract states $\overline{S}$ and abstract actions $\overline{A}$, a Generalized Policy Automaton (GPA) $\overline{G} = \langle \overline{V}, \overline{E} \rangle$ is a non-deterministic finite-state automaton that can be represented as an abstract directed hypergraph where $\overline{V} = \overline{S}$ is the vertex set and $\overline{E} \subseteq \overline{V} \times \mathcal{P}(\overline{V}) \times \overline{A}$ is the set of directed hyperedges s.t. each hyperedge $\overline{e} \in \overline{E}$ is a tuple $(\overline{e}_{start}, \overline{e}_{results}, \overline{e}_{action})$ representing a start vertex, a set of result vertices, and an action label.

3.2 Canonical Abstraction

A key ingredient in learning generalized knowledge using feature-based abstractions is using *rich* feature sets that facilitate the learning of compact generalized policies. We use canonical abstractions for automatically synthesizing a set of such rich, domain-independent features.

Canonical abstractions, commonly used in program analysis, have been shown to be useful in generalized planning (Srivastava et al., 2011; Karia and Srivastava, 2021). Given a predefined set of abstraction predicates, canonical abstractions group together different objects based on the subset of abstraction predicates that they satisfy in a state. Each subset of abstraction predicates is known as a *summary* element. We used the set of all unary predicates as abstraction predicates in this paper.[1] We use these abstraction predicates to obtain a set of features as described later in the paper.

Let ψ be a summary element, then, we define $\phi_\psi(s)$ as a function that returns the set of objects that satisfy ψ in a concrete state s. An object o belongs to exactly one summary element in any given concrete state s and this is the *maximal* summary element for that object. Similarly, for any given predicate $p_n \in \mathcal{P}$ where n is the arity, $\phi_{p_n}(\psi_1, \ldots, \psi_n)$ is defined as the set of all n-ary predicates in s that are consistent with the summary elements composing the predicate $p_n(\psi_1, \ldots, \psi_n)$, i.e., $\phi_{p_n(\psi_1,\ldots,\psi_n)}(s) = \{p_n(o_1, \ldots, o_n) | p_n(o_1, \ldots, o_n) \in s, o_i \in \phi_{\psi_i(s)}\}$.

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

Let Ψ be the set of all summary elements and $\mathcal{P}_i$ be the set of all predicates $p \in \mathcal{P}$ 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 summary elements. 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 $\overline{s}$ as a total valuation of F_α 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 $a(o_1, \ldots, o_n)$ when applied to s returns an abstract action $\overline{a}(\psi_1, \ldots, \psi_n)$ where $\overline{a} \equiv a$ and ψ_i is the summary element that object o_i satisfies, i.e., $o_i \in \phi_{\psi_i}(s)$ for $\psi_i \in \Psi$.

[1] 0-ary predicates are represented as unary predicates with a default "phantom" object.

Algorithm 1 GPA-accelerated SSP Solver

Require: SSP $P = \langle O, S, A, s_0, g, T, C \rangle$, GPA $\overline{G} = \langle \overline{V}, \overline{E} \rangle$, Feature Sets F_α, F_β

1: $C' = C$ {copy over the cost function of P}
2: **for** $(s, a, s') \in S \times A \times S$ **do**
3: $\quad \overline{s}, \overline{a}, \overline{s'} \leftarrow \alpha(F_\alpha, s), \beta(F_\beta, s, a), \alpha(F_\alpha, s')$
4: $\quad$ **if** there is no edge $\overline{e} \in \overline{E}$ s.t. $\overline{e}_{start} = \overline{s}, \overline{s'} \in \overline{e}_{results}$, and $\overline{e}_{action} = \overline{a}$ **then**
5: $\quad\quad C'[s, a, s'] = \infty$
6: $\quad$ **end if**
7: **end for**
8: $P|_{\overline{G}} = \langle O, S, A, s_0, g, T, C' \rangle$
9: $v_{P|_{\overline{G}}}, \pi_{P|_{\overline{G}}} \leftarrow$ SSP solver$(P|_{\overline{G}})$
10: **if** $\pi_{P|_{\overline{G}}}$ is a partial proper policy **then**
11: $\quad$ **return** $\pi_{P|_{\overline{G}}}$
12: **else**
13: $\quad$ Initialize v_P using $v_{P|_{\overline{G}}}$
14: $\quad v_P, \pi_P \leftarrow$ SSP solver(P)
15: $\quad$ **return** π_P
16: **end if**

Example Consider the state s_{eg} from the running example (Sec. 2). Let the summary elements computed using canonical abstraction be $\psi_1 = \{\textit{location}\}$, $\psi_2 = \{\textit{rock}\}$, and $\psi_3 = \{\textit{in-rover}, \textit{rock}\}$ and let $\Psi_{eg} = \{\psi_1, \psi_2, \psi_3\}$ be the set of all summary elements. We can compute $\phi_{\psi_1}(s_{eg}) = \{\textit{base}, l_1, l_2\}$, $\phi_{\psi_2}(s_{eg}) = \{r_1\}$, and $\phi_{\psi_3}(s_{eg}) = \{r_2\}$. Similarly, $\phi_{\textit{rock-at}(\psi_1,\psi_1)} = \{\}$, $\phi_{\textit{rock-at}(\psi_1,\psi_2)} = \{\}$, $\phi_{\textit{rock-at}(\psi_2,\psi_1)} = \{\textit{rock-at}(r_1, l_2)\}$, and so on. The abstract state $\overline{s}_{eg} = \{\psi_1 = 2, \psi_2 = 1, \psi_3 = 1, \textit{rock-at}(\psi_1, \psi_1) = 0, \textit{rock-at}(\psi_1, \psi_2) = 0, \textit{rock-at}(\psi_2, \psi_1) = \frac{1}{2}, \ldots\}$. Similarly, for an action, unload$(r_2, \textit{base})$, in s_{eg}, the abstract action would be unload(ψ_3, ψ_1).

3.3 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) and compute GPAs iteratively from a small training set containing solutions of similar SSP instances as outlined below.

To form our training set, we use a library of solution policies $\Pi = \{\pi_1, \ldots, \pi_n\}$ for *small* problems $P_1, \ldots, P_n$ that can be easily computed by existing SOA SSP solvers. We represent each policy $\pi_i \in \Pi$ as a concrete directed hypergraph and construct a training set $\mathcal{T} = \{G_{\pi_1}, \ldots, G_{\pi_n}\}$. We initialize an initial empty GPA $\overline{G} = \langle \{\overline{V}, \overline{E}\} \rangle = \langle \{\}, \{\} \rangle$. Next, for each $G_{\pi_i} = \langle V_{\pi_i}, E_{\pi_i} \rangle \in \mathcal{T}$, we merge the concrete directed hypergraph with $\overline{G}$ as follows: $\overline{V} = \overline{V} \cup \{\alpha(F_\alpha, v) | v \in V_{\pi_i}\}$. Next, we form a set of abstract hyperedges by converting every edge in E_{π_i} to its corresponding abstract counterpart, i.e., $\overline{E}_{\pi_i} = \{\langle \overline{e}_{start}, \overline{e}_{results}, \overline{e}_{action} \rangle\}$ for every edge $e \in E_{\pi_i}$ s.t. $\overline{e}_{start} = \alpha(F_\alpha, e_{start})$, $\overline{e}_{results} = \{\alpha(F_\alpha, v) | v \in e_{results}\}$ and $\overline{e}_{action} = \beta(F_\beta, e_{start}, e_{action})$. Finally, we merge this set of abstract hyperedges using the following procedure: $\overline{E} = \overline{E} \cup \overline{E}_{\pi_i}$. Once this is done, we compress the GPA by replacing any edges in $\overline{E}$ that have the same start nodes and labels with a new edge, i.e., for any two edges $\overline{e}^1, \overline{e}^2 \in \overline{E}$ s.t. $\overline{e}^1_{start} = \overline{e}^2_{start}$, and $\overline{e}^1_{action} = \overline{e}^2_{action}$, we delete the two edges and replace the edge set with a new merged edge: $\overline{E} = \langle \overline{e}^1_{start}, \overline{e}^1_{results} \cup \overline{e}^2_{results}, \overline{e}^1_{action} \rangle \cup \overline{E} \setminus \{\overline{e}^1, \overline{e}^2\}$.

3.4 Solving SSPs using GPAs

Given a GPA $\overline{G}$, we now describe our method for solving SSPs. To do so, we modify the cost function of an SSP P to generate a new SSP $P|_{\overline{G}}$, taking into account the transitions in $\overline{G}$ as shown in Alg. 1 (lines 1-7). Our SSP solver -- GPA-accelerated SSP Solver -- operates as follows: Create a copy of the cost function of the original SSP P (line 1); Iterate over each (s, a, s') tuple of P and convert them to $(\overline{s}, \overline{a}, \overline{s'})$ tuples by computing abstract states and action representations using canonical abstraction (lines 2-3); Check if a hyperedge corresponding to the tuple $(\overline{s}, \overline{a}, \overline{s'})$ exists in the GPA $\overline{G}$. A hyperedge $(e_{start}, e_{results}, e_{action}) \in \overline{E}$ is corresponding to the tuple $(\overline{s}, \overline{a}, \overline{s'})$ iff

$e_{start} = \overline{s}$, $\overline{s'} \in e_{results}$, and $e_{action} = \overline{a}$ (line 4); If the hyperedge does not exist in the GPA $\overline{G}$, we set the cost entry for the concrete tuple (s, a, s') to infinity (line 5). Finally, we create a GPA-constrained SSP $P|_{\overline{G}}$ that is same as the original SSP except for the cost function which is a modified copy of the original cost function (line 8), and use any off-the-shelf SSP solver that guarantees to find a partial proper policy if one exists to solve $P|_{\overline{G}}$ (line 9).

The goal of modifying the cost function is to prevent transitions whose abstract translations are not present in the Generalized Policy Automaton $\overline{G}$ to be used when performing Bellman updates for the new SSP $P|_{\overline{G}}$. As a result, actions belonging to such transitions cannot appear in $\pi_{P|_{\overline{G}}}$. As a consequence, the existence of a partial proper policy is not guaranteed in $P|_{\overline{G}}$. Nevertheless, the GPA-accelerated SSP solver is guaranteed to return a partial proper policy. This is because, if the computed policy for $P|_{\overline{G}}$ is a partial proper policy, the GPA-accelerated SSP solver directly returns it (lines 10-11). Otherwise, if a partial proper policy for $P|_{\overline{G}}$ does not exist or an improved policy is desired, then the GPA-accelerated SSP solver re-uses the computational effort expended in $P|_{\overline{G}}$ to initialize the value function of P using $v_{P|_{\overline{G}}}$ (line 13) and uses the SSP solver to compute a policy for the original SSP P (lines 14-15). This policy is guaranteed to be a partial proper policy if the SSP solver used is complete (Thm. 3.2). Also, if the used SSP solver is optimal then the bootstrapped policy computed for P is at least as cost-effective as the policy computed for $P|_{\overline{G}}$ (Thm. 3.1).

The GPA-accelerated SSP solver computes a policy for $P|_{\overline{G}}$ in the space of cross-product of the states of the GPA $\overline{G}$ with the states of the SSP P, similar to HAMs (Parr and Russell, 1997). This policy, defined as a *hierarchically optimal policy*, corresponds to an optimal policy for P among all the policies that satisfy the constraints encoded by the GPA $\overline{G}$. As seen in our empirical analysis in Sec. 4, we observe that using a small set of example policies that are sufficient to capture rich generalized control structures results in GPA-accelerated SSP solver returning partial proper policies within a fraction of the original computational effort. We now prove theoretical guarantees of completeness and hierarchical optimality below.[2]

Theorem 3.1. Given an SSP P, an optimal SSP solver, a GPA $\overline{G}$, and the corresponding GPA-constrained SSP $P|_{\overline{G}}$, let v_P, $v_{P|_{\overline{G}}}$, and π_P, $\pi_{P|_{\overline{G}}}$ be the state-value functions and policies obtained using Alg. 1 respectively, then π_P is at least as cost-effective as $\pi_{P|_{\overline{G}}}$, i.e., $v_P(s_0) \leq v_{P|_{\overline{G}}}(s_0)$.

Proof (Sketch). Our proof is based on the following intuition. In the trivial case where $\pi_{P|_{\overline{G}}}$ is not a partial proper policy (line 10), $v_{P|_{\overline{G}}}(s_0) = \infty$ and solving P using bootstrapped estimates (line 13) is equivalent to solving P without bootstrapping. If $\pi_{P|_{\overline{G}}}$ is a partial proper policy, then it is also a partial proper policy for P since Alg. 1 does not change the transition function of P and thus $v_P(s_0) = v_{P|_{\overline{G}}}(s_0)$ when bootstrapping. An optimal solver using the Bellman equations will only accept a new policy π over $\pi_{P|_{\overline{G}}}$ if it leads to improvement over the *current, greedy* policy i.e., $\pi_{P|_{\overline{G}}}$. Thus, the only case in which Alg. 1 would return a different policy is if it finds a partial proper policy whose expected cost is less than or equal to that of $\pi_{P|_{\overline{G}}}$ i.e., $v_P(s_0) \leq v_{P|_{\overline{G}}}(s_0)$. □

Theorem 3.2. Given a training set $\mathcal{T} = \{G_{\pi_1}, ..., G_{\pi_n}\}$ containing partial proper policies, an optimal, complete SSP solver, feature sets F_α, F_β, a learned GPA $\overline{G}$, and an SSP P, the GPA-accelerated SSP Solver (Alg. 1) is guaranteed to be complete, i.e., return a partial proper policy for P if one exists.

Proof (Sketch). Line 10 of the Alg. 1 checks if the policy computed on solving $P|_{\overline{G}}$ is a partial proper policy and returns it in case it is partial proper policy. If the policy is not a partial proper policy then it uses the bootstrapped value estimates and solves P. If the used SSP solver is complete, then Alg. 1 is guaranteed to return a partial proper policy if one exists. □

Theorem 3.3. Given a training set $\mathcal{T} = \{G_{\pi_1}, ..., G_{\pi_n}\}$ containing partial proper policies, an optimal complete SSP solver, feature sets F_α, F_β, a learned GPA $\overline{G}$, and an SSP P, the GPA-accelerated SSP Solver (Alg. 1) is guaranteed to return a hierarchically optimal policy for P if one exists.

Proof (Sketch). We provide a proof by contradiction. Suppose that an optimal policy $\pi_* \neq \pi_{P|_{\overline{G}}}$ exists for $P|_{\overline{G}}$. This implies that for at least one of the states s in $P|_{\overline{G}}$, $v_*(s) < v_{\pi_{P|_{\overline{G}}}}(s)$ and as a result $\pi_*(s) \neq \pi_{P|_{\overline{G}}}(s)$. This implies the existence of an optimal action a for which the cost function

[2]Complete proofs for all theorems are available in the supplementary material.

$C'[s, a, s']$ was set to ∞ since otherwise this action a would have been considered by the SSP solver. This implies that a transition (s, a, s') in the training data was not captured by the GPA and is a contradiction since no transitions are discarded when creating $\overline{G}$ from the training data $\mathcal{T}$. $\square$

4 Experiments

We conducted an empirical evaluation on five well-known benchmark domains that were selected from the International Planning Competition (IPC) and International Probabilistic Planning Competition (IPPC) (Younes et al., 2005) as well as robotic planning (Shah et al., 2020). As a part of our analysis, we aim to determine if 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 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. Also, modern SOA solvers for PPDDL are available.

For our baselines, we focus on complete solvers for SSPs. We used Labeled RTDP (LRTDP) (Bonet and Geffner, 2003) and Soft-FLARES (Pineda and Zilberstein, 2019) which are state-of-art (SOA), complete SSP solvers. These algorithms internally generate their own heuristics using the domain and problem file. We conducted an empirical evaluation and used the FF heuristic (Hoffmann, 2001) as the internal heuristic function for all algorithms. We used FF even though it is inadmissible 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 Pineda and Zilberstein (2019) to Python.[3] We utilized problem generators from the IPPC suite 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 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 balls b initially located in room A to room B. We modified the gripper to be slippery so that picking balls have 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 the two locations preferring one location with probability 0.6. Despite this simple setting, the keva domain has been shown to be a challenging robotics problem (Shah et al., 2020).

Delicate Can(c) An arrangement of c cans on a table of which 1 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 our training set, we used at most ten optimal 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. This highlights the advantages of GPAs that can quickly be learned in a few-shot setting.

Test Setup We fixed the time and memory limit for each problem to 7200 seconds and 16GiB 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

[3] Our source code, training and test datasets, and instructions to reproduce the reported results is included in the supplementary material.

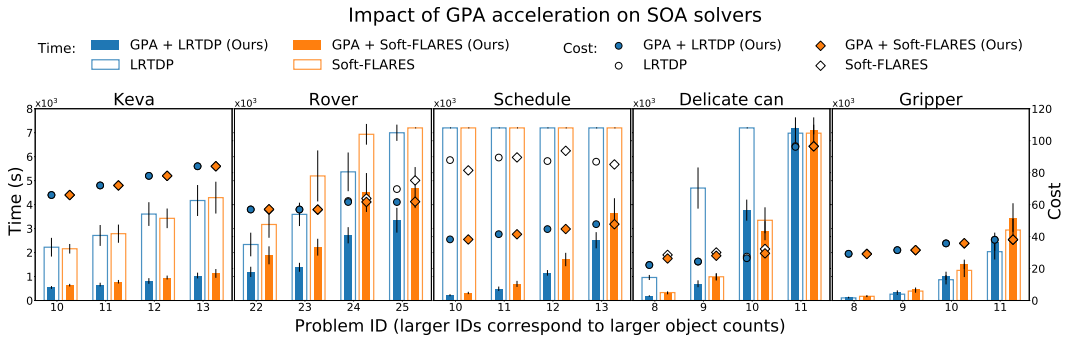

Figure 1: Impact of learned GPAs on solver performance (lower values are better). Left y-axes and bars show solution times (in units of 1000 secs) for our approach and baseline SOA solvers (LRTDP and Soft-FLARES). Right y-axes and points show cost incurred by the policy computed by our approach and the 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. This and additional empirical data is available in the supplementary material.

and test problems and their parameters, used hyperparameters for configuring baselines, etc., are included in the supplementary material.

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 1 standard deviation. Results of our experiments are illustrated in Fig. 1.

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 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 GPA to still 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 a domain expert after employing significant effort.

Finally, because of the fixed timeout used, the maximum time of the baselines was bounded, making the impact of GPA 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 ϵ-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 fail to learn any generalized knowledge and as result cannot readily scale to larger problems with a greater number of objects.

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.

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 (Bonet et al., 2019; 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 or its cost and can only work on deterministic problems.

Our approach differs from these approaches in several aspects. Our approach constructs a GPA automatically without any human intervention. Using canonical abstraction, we lift problem-specific characteristics like object names and object counts. Another key difference between other techniques is that our approach can easily incorporate solutions from new examples into the controller without having to remember any of the earlier examples. This allows our learning to scale better and can naturally utilize *leapfrogging* (Groshev et al., 2018a; Karia and Srivastava, 2021) 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.

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 allowing our approach to compute policies of comparable cost in a fraction of the effort. Our approach comes with guarantees of hierarchical optimality given the training data and also comes with the guarantee of always finding a solution provided one exists.

There are several interesting research directions for future work. Currently, GPAs do not employ any memory and can be improved by incorporating a finite amount of memory for better pruning of the action transitions. Description Logics (DL) are more expressive than canonical abstractions and have been demonstrated by Bonet et al. (2019) to be effective at synthesizing memoryless controllers for deterministic planning problems. Our approach can easily utilize any relational abstraction and it would be interesting to evaluate the efficacy of description logics as compared to canonical abstractions. 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 Gripper. We plan to investigate these directions of research in future work.

References

J. S. Aguas, S. J. Celorrio, and A. Jonsson. Hierarchical finite state controllers for generalized planning. In *IJCAI*, 2016.

A. Bai and S. J. Russell. Efficient reinforcement learning with hierarchies of machines by leveraging internal transitions. In *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 *ICAPS*, 2003.

B. Bonet, H. Palacios, and H. Geffner. Automatic derivation of memoryless policies and finite-state controllers using classical planners. In *ICAPS*, 2009.

B. Bonet, G. Francès, and H. Geffner. Learning features and abstract actions for computing generalized plans. In *AAAI*, 2019.

C. Boutilier, R. Reiter, and B. Price. Symbolic dynamic programming for first-order mdps. In *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 *AAAI*, 2021.

S. Garg, A. Bajpai, and Mausam. Symbolic network: Generalized neural policies for relational mdps. In *ICML*, 2020.

E. Groshev, M. Goldstein, A. Tamar, S. Srivastava, and P. Abbeel. Learning generalized reactive policies using deep neural networks. In *ICAPS*, 2018a.

E. Groshev, A. Tamar, M. Goldstein, S. Srivastava, and P. Abbeel. Learning generalized reactive policies using deep neural networks. In *AAAI Spring Symposium Series*, 2018b.

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 *AAAI*, 2021.

A. Kolobov, Mausam, and D. S. Weld. A theory of goal-oriented mdps with dead ends. In *UAI*, 2012.

A. Levy, G. Konidaris, R. Platt, and K. Saenko. Learning multi-level hierarchies with hindsight. In *ICLR*, 2019.

R. Parr and S. J. Russell. Reinforcement learning with hierarchies of machines. In *NeurIPS*, 1997.

L. E. Pineda and S. Zilberstein. Soft labeling in stochastic shortest path problems. In *AAMAS*, 2019.

V. Pong, S. Gu, M. Dalal, and S. Levine. Temporal difference models: Model-free deep rl for model-based control. In *ICLR*, 2018.

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`, 2010.

S. Sanner and C. Boutilier. Approximate linear programming for first-order mdps. In *UAI*, 2005.

N. Shah, D. K. Vasudevan, K. Kumar, P. Kamojjhala, and S. Srivastava. Anytime integrated task and motion policies for stochastic environments. In *ICRA*, 2020.

S. Srivastava, N. Immerman, and S. Zilberstein. A new representation and associated algorithms for generalized planning. *Artif. Intell.*, 175(2):615–647, 2011.

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 NeurIPS*, 29, 2016.

S. Toyer, F. W. Trevizan, S. Thiébaux, and L. Xie. Action schema networks: Generalised policies with deep learning. In *AAAI*, 2018.

F. W. Trevizan and M. M. Veloso. Short-sighted stochastic shortest path problems. In *ICAPS*, 2012.

J. Wu and R. Givan. Discovering relational domain features for probabilistic planning. In *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.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Sec. 6
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] Proof sketches are included in the main paper and complete proofs are included in the supplementary material as mentioned in footnote 2.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Source code as well as detailed instructions to reproduce our experiments are included in the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] The main paper describes significant details of training. Detailed descriptions of the entire setup is included in the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Fig. 1 shows error bars for the times. Standard deviations for costs are included in the supplementary material as is noted in the caption for Fig. 1.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Sec. 4
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We appropriately cite the problem generators that we used in Sec. 4.
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]