Generating Data In Planning: SAS$^+$ Planning Tasks of a Given Causal Structure

Abstract

The need for data in planning has long been established, by, e.g., machine learning based approaches. The existing data, however, is quite limited. There exists only a relatively small amount of hand-crafted planning domains, mostly introduced through International Planning Competitions. Further, this collection of domains is not necessarily diverse: many of these domains are some variants of the transportation problem.

In this work, we alleviate the shortage in existing planning tasks by automatically generating tasks of a particular causal structure. Given any graph G, we show how to create a SAS$^+$ planning task with the causal graph isomorphic to G. We create a large collection of planning tasks by randomly generating graphs of various structural restrictions and creating planning tasks for these graphs, but also, more importantly, we provide the community with a tool that allows for on-demand generation of additional, possibly larger tasks. Our experimental evaluation ensures that the generated collection is interesting for the current state of affairs in classical cost-optimal planning, showing the performance of state-of-the-art symbolic search and heuristic search based planners.

Introduction

Since the first planning tasks encoded in STRIPS language back in 1971 (Fikes and Nilsson 1971), data, a.k.a. planning tasks, was the corner stone and one of the main drivers of research in planning. With the beginning of International Planning Competitions (IPC) in 1998 (McDermott 2000) came the increase in the availability of planning tasks, with the current estimate of slightly over 70 domains, including some variants in different formalisms. All these domains are hand-crafted, although some correspond to machine translation from a different problem (Palacios and Geffner 2009; Bonet, Palacios, and Geffner 2009; Grastien and Scala 2018; Sohrabi et al. 2018). Not only that most of these domains are hand-crafted, the collection is not necessarily diverse. Many of these domains are some variants of the transportation problem.

A major focus in classical planning was on heuristic search, with heuristics automatically obtained for planning tasks, exploiting the task structure. Few examples that explicitly exploit the causal structure include the *causal graph heuristic* (Helmert 2004) and the *structural pattern heuristics* (Katz and Domshlak 2010; Katz and Keyder 2012). Others, such as *pattern databases* (PDBs) (Edelkamp 2001) exploit the causal information in e.g., pattern selection (Haslum et al. 2007). *Merge-and-shrink heuristics* (Helmert, Haslum, and Hoffmann 2007) use the causal graph for guiding the merge process. Most existing heuristics that work on the multi-valued representation exploit the causal information in one way or another. Further, starting with the seminal work of Bäckström and Nebel (1995), the research on the complexity of planning tasks had a major focus on the characterization of planning fragments by their causal graph structure (Domshlak and Brafman 2002; Katz and Domshlak 2007; 2008; Giménez and Jonsson 2008; 2009; Katz and Keyder 2012), showing these fragments to belong to a variety of complexity classes. For these two reasons, various planners performance heavily relies on the various structural characteristics of the input planning task.

The aim of this work is to generate planning tasks of a specific predefined structure. Here, we focus on the characterization of planning tasks by the structure of their causal graphs. Given a collection of multi-valued variables and a graph representing causal connections between these variables, we propose a way of generating SAS$^+$ actions, initial state and a goal, in a way that the causal graph of the resulting task will match the input graph. Our aim is to be able to automatically generate a diverse collection of planning tasks, as large as needed for various purposes. One such example purpose is learning a good planner selection strategy (Sievers et al. 2019; Ma et al. 2020). Another possible purpose is an additional source of benchmarks for empirical evaluation of new planning algorithms. We test the generated collection with two modern cost-optimal planners that represent the two popular state-of-the-art approaches to cost-optimal planning. For symbolic search, we chose the planner SymBA* (Torralba et al. 2014), winner of the sequential optimal track of International Planning Competition 2014 and one of the planners in the winning portfolio of IPC 2018 (Katz et al. 2018). For heuristic search, we chose A^* with LM-cut heuristic (Helmert and Domshlak 2009), a compo-

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

nent of many modern heuristics search based planners. Our experiments confirm that (i) the generated collection is challenging for both heuristic search and symbolic search based planners, and (ii) there is no clear dominance to any of the techniques.

The rest of the paper is structured as follows. We start with introducing the planning formalism and the notation used throughout the paper. We then move to construction, where we first describe various causal graph structures and the way these graphs can be constructed, and then describe the construction of planning tasks given a causal graph. Next, we present the experimental evaluation, including describing the way we have created our collection. Finally, we discuss the related work, and conclude with the summary of our results and future work.

Preliminaries

A SAS$^+$ *planning task* (Bäckström and Nebel 1995) is given by a tuple $\langle \mathcal{V}, A, s_0, s_* \rangle$, where $\mathcal{V}$ is a set of *state variables*, A is a finite set of *actions*. Each state variable $v \in \mathcal{V}$ has a finite domain $dom(v)$. Each pair $\langle v, \vartheta \rangle$ of variable $v \in \mathcal{V}$ and its value $\vartheta \in dom(v)$ is called a *fact*. By F_v we denote the set $\{\langle v, \vartheta \rangle \mid \vartheta \in dom(v)\}$ of facts for the variable v, and the set of all facts is denoted by $F := \bigcup_{v \in \mathcal{V}} F_v$. A (partial) assignment to the variables $\mathcal{V}$ is called a *(partial) state*. Often it is convenient to view partial state p as a set of facts with $\langle v, \vartheta \rangle \in p$ if and only if $p[v] = \vartheta$. For a partial assignment p, $\mathcal{V}(p) \subseteq \mathcal{V}$ denotes the subset of state variables instantiated by p. Partial state p is *consistent* with state s if $p \subseteq s$. We denote the set of states of a planning task by S. s_0 is the *initial state*, and the partial state s_* is the *goal*. Each *action* a is a pair $\langle pre(a), \mathit{eff}(a) \rangle$ of partial states called *preconditions* and *effects*. An *action cost* is a mapping $C : A \to \mathbb{R}^{0+}$. An action a is applicable in a state $s \in S$ if and only if $pre(a)$ is consistent with s. Applying a changes the value of v to $\mathit{eff}(a)[v]$, if defined. The resulting state is denoted by $s[\![a]\!]$. An action sequence $\pi = \langle a_1, \ldots, a_k \rangle$ is applicable in s if there exist states $s_0, \cdots, s_k$ such that (i) $s_0 = s$, and (ii) for each $1 \leq i \leq k$, a_i is applicable in s_{i-1} and $s_i = s_{i-1}[\![a_i]\!]$. We denote the state s_k by $s[\![\pi]\!]$. π is a plan iff π is applicable in s_0 and s_* is consistent with $s_0[\![\pi]\!]$. We denote by $\mathcal{P}(\Pi)$ (or just $\mathcal{P}$ when the task is clear from the context) the set of all plans of Π. The cost of a plan π, denoted by $C(\pi)$ is the summed cost of the actions in the plan.

A central role in what follows is played by a standard structure in classical planning, called *causal graph* (Helmert 2004). The causal graph CG_Π of a task Π is a digraph with vertices $\mathcal{V}$. An arc (v, v') is in CG_Π iff $v \neq v'$ and there exists an action $a \in A$ such that $(v, v') \in [\mathcal{V}(\mathit{eff}(a)) \cup \mathcal{V}(pre(a))] \times \mathcal{V}(\mathit{eff}(a))$. For an action a, by E_a we denote the set of all such arcs, and by $E_{A'}$ we denote the union of all sets of arcs E_a for $a \in A'$.

Another structure typically used in planning for computing relaxation based heuristics is *relaxed planning graph* (Hoffmann and Nebel 2001), which is a layered graph of facts and actions, describing action application in the planning task, under value accumulating semantic. The layers are added until a fixpoint is reached, that is no new fact can be achieved. The first fact layer F_1 thus corresponds to the facts from the initial state, and the last layer is also a fact layer, and it is equal to the preceding fact layer. Each action layer A_i consists of all actions from A that are applicable in F_i, that is $A_i = \{a \in A \mid pre(a) \subseteq F_i\}$. The next fact layer F_{i+1} is then constructed by adding to F_i all facts achieved by the actions in A_i, namely $F_{i+1} = F_i \cup \bigcup_{a \in A_i} \mathit{eff}(a)$.

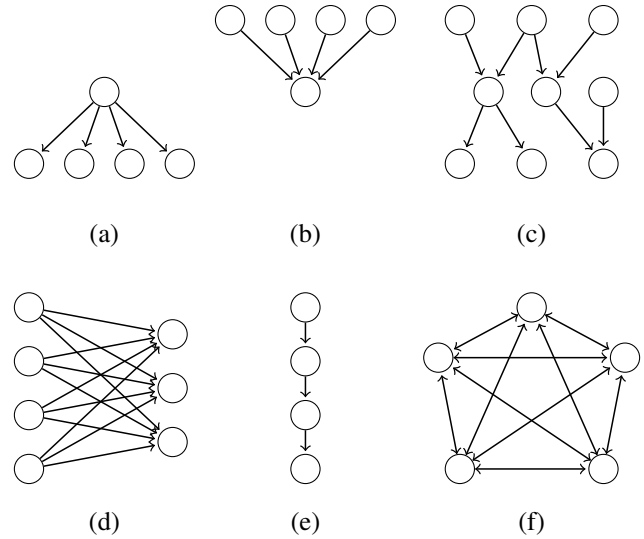

Figure 1: Selected causal graph structures: (a) fork, (b) inverted fork, (c) polytree, (d) directed bipartite graph, (e) chain, and (f) complete graph .

Finally, SAS$^+$ representation is often not provided directly, and is *translated* from STRIPS representation (Helmert 2006). The multi-valued variables in SAS$^+$ then correspond to invariant groups of pairwise mutually exclusive facts (mutexes), where exactly one such fact is true in any state reachable from the initial state. Each such invariant group over STRIPS facts corresponds to a set of facts *at most one* of which can be true in any given state that is reachable from the initial state. If there exist such states where no facts are true, then an additional value is added, representing that none of the facts in the invariant group is true.

Construction

We start our construction by defining a graph to be served as the causal graph of the constructed task. For that, we focus here on the following causal graph structures: *chain*, *directed chain*, *fork*, *inverted fork*, *bipartite graph*, *directed bipartite graph*, *tree*, *polytree*, and *directed acyclic graph*, *complete graph*, and *random graph*. For some of these structures, namely *directed chain*, *fork*, *inverted fork*, and *complete graph* the graphs are fully defined by the number of nodes (modulo automorphisms). In other cases, we introduce randomness into the graph construction. In what follows, we first describe how we handle these cases, and then how a task with the given causal structure is constructed.

Directed Bipartite Graph: A full directed bipartite graph is constructed by first randomly partitioning the nodes

into left and right and then introducing an edge from each node on the left to each node on the right.

Bipartite Graph: A full undirected bipartite graph is constructed by first randomly partitioning the nodes into left and right and then introducing an edge from each node on the left to each node on the right, and vice versa.

Chain: The chain of n nodes $v_1, \ldots, v_n$ is created as follows. For each $1 \leq i < n$, we randomly decide whether to add an edge (v_i, v_{i+1}), with probability p. If no such edge is added, we add the edge (v_{i+1}, v_i) and if the edge (v_i, v_{i+1}) was added, we decide with probability p whether to add the edge (v_{i+1}, v_i).

Tree: A directed tree of n nodes $v_1, \ldots, v_n$ is constructed by choosing for each node v_i a parent randomly out of the nodes $v_1, \ldots, v_{i-1}$.

Polytree: For a polytree, we start with a tree constructed as above, and then for each edge switch its direction with probability p.

Directed Acyclic Graph: A directed acyclic graph of n nodes $v_1, \ldots, v_n$ is constructed by choosing for each node v_i at least one parent randomly out of the nodes $v_1, \ldots, v_{i-1}$. We do that by going over all the preceding nodes and deciding with probability p whether to add an edge from the preceding node to the current node. If no edges were added, we repeat until at least one edge is added for each node (except the first one).

Random Graph: For each pair of nodes v_i and v_j we randomly decide whether to add a directed edge from v_i to v_j.

Figure 1 exemplifies selected graph structures. Having described how a causal graph for the future planning task is constructed, we now switch to the next step, showing how to construct a planning task with that causal graph.

Planning Task Construction

Given a graph $G = (\mathcal{V}, E)$ and a number of facts $n \geq 2|\mathcal{V}|$, we construct the SAS$^+$ planning task $\Pi = \langle \mathcal{V}, A, s_0, s_* \rangle$ with the causal graph G as follows. First, we choose the domain size $d_v \geq 2$ for each of the multi-valued variables $v \in \mathcal{V}$ and assume w.l.o.g. the values to be $dom(v) = \{0, \ldots, d_v - 1\}$. The variables represent sets of mutually exclusive facts (mutexes), and each such set corresponds to one of the two types of mutexes, namely, either exactly one or at most one of the values is true in all reachable states. We randomly decide which variables belong to which category. For the variables that represent the at-most-one case, we dedicate the last domain value to represent the case when none of the other facts are true. Next, w.l.o.g. we assume $s_0[v] = 0$ for all $v \in \mathcal{V}$. Then, we construct the actions, in layers, while constructing the relaxed planning graph. Finally, the goal is chosen from the last fact layer of the relaxed planning graph, making sure that at least one of the chosen facts is unique to the last fact layer and that at most one fact is chosen per variable. In what follows, we describe how actions are constructed. Starting with the initial state as the first fact layer F_0, we create an action layer L_i by

Algorithm 1 Construction of a planning task according to a given causal graph structure.

Input: Graph $G = (\mathcal{V}, E)$, number of facts $n \geq 2|\mathcal{V}|$

Partition n into $|\mathcal{V}|$ values $d_v \geq 2$ such that $\sum_{v \in \mathcal{V}} d_v = n$
$F_v \leftarrow \{\langle v, \vartheta \rangle \mid 0 \leq \vartheta < d_v\}$ for all $v \in \mathcal{V}$
$s_0[v] \leftarrow 0$ for all $v \in \mathcal{V}$
$k \leftarrow 0$
$A \leftarrow \emptyset$
$F_k \leftarrow s_0$
while $|F_k| < n$ **or** $E \setminus E_A \neq \emptyset$ **do**
  $k \leftarrow k + 1$
  $m_k \leftarrow$ number of new facts for layer k
  $F_k, A_{k-1} \leftarrow$ CREATELAYER(m_k, F_{k-1}, A)
  $A \leftarrow A \cup A_{k-1}$
Select $s_* \subseteq F_k$ such that $\forall v \in \mathcal{V}, |s_* \cap F_v| \leq 1$ and $s_* \cap (F_k \setminus F_{k-1}) \neq \emptyset$
return $\Pi = \langle \mathcal{V}, A, s_0, s_* \rangle$

function CREATELAYER(m, F, A)
  $A' \leftarrow \emptyset$
  $F' \leftarrow F$
  while $|F' \setminus F| < m$ **or** ($m = 0$ **and** $E \setminus E_{A \cup A'} \neq \emptyset$) **do**
    $a \leftarrow$ CREATEACTION
    if $E_a \subseteq E$ **then**
      if $E_a \setminus E_{A \cup A'} = \emptyset$ **and** Random(p) **then**
        Continue
      $A' \leftarrow A' \cup \{a\}$
      $F' \leftarrow F' \cup \mathit{eff}(a)$
  return F', A'

(I) selecting a subset of facts from the fact layer F_i, ensuring at most one fact is selected per variable,

(II) partitioning the selected set of facts into prevail condition and precondition, and

(III) choosing for all[1] the variables of the precondition facts a different value as its effect.

The constructed action is checked against the graph G, ensuring that it contributes only edges that exist in G. If not, the action is discarded. Additionally, if the constructed action does not add any new causal edges and does not achieve new facts, we randomly decide whether to keep it.

The generic approach to action construction described above can be adapted to enforce particular properties. We discuss three such cases in detail.

(A) The first case is enforcing the action to achieve at least one new fact. For that, one can ensure in steps (I) and (II) that the precondition includes facts for some variables $v \in \mathcal{V}$ that are not fully covered by the fact layer F_i (that is $F_v \setminus F_i \neq \emptyset$), and in step (III) to choose one

[1] While SAS$^+$ representation does not require to specify the precondition when the effect is specified, in order to ensure maintaining variables as mutexes of facts, we restrict ourselves here to always specifying the precondition in such cases.

of these facts $F_v \setminus F_i$. Note that this can be done without adding any edges to the causal graph, if a single fact is chosen in step (I).

(B) The second case is enforcing the preconditions to include atoms from $F_i \setminus F_{i-1}$, enforced in step (I).

(C) The third case is enforcing adding a particular edge $\langle v, v' \rangle$ to the causal graph. This can be done by ensuring that $\langle v, \vartheta \rangle$ and $\langle v', \vartheta' \rangle$ are chosen in step (I), and in step (II) at most one of these facts is chosen for the prevail condition.

We randomly independently decide whether to enforce the options (A)-(C) and whether to add edges to the causal graph. Note that not all combinations are always possible. In such cases, an action is not constructed in that iteration. Each layer is constructed until a sufficient number of new facts $F_{i+1} \setminus F_i$ is added. The construction is stopped when all facts were achieved and all edges from G are reflected in the causal graph of the constructed planning task. The latter is enforced in the last layer. Algorithm 1 describes the construction of a planning task from the given graph G, where the function CREATEACTION creates a single action, randomly choosing among the options described above.

Theorem 1 *Given G and n, Algorithm 1 terminates in time polynomial in $|G|$ and n and returns a planning task with the causal graph G.*

Proof: The proof follows from the fact that in line 18 of Algorithm 1, for some of the options for action creation must eventually hold $E_a \subseteq E$ and $E_a \setminus E_{A \cup A'} \neq \emptyset$. Therefore, CREATELAYER terminates and returns a layer with m new facts or, if $m = 0$, with A' such that $E_{A \cup A'} = E$. As there are only a constant number of options, a new fact is achieved or a new causal graph edge is covered in time $O(1)$ and therefore CREATELAYER terminates in time $O(m + |E|)$. Since at least one new fact is added in each layer, Algorithm 1 terminates in time $O(n|E|)$. Since the while loop in line 7 terminates only when $E \setminus E_A = \emptyset$ and actions a are added to A only if $E_a \subseteq E$, when the algorithm terminates we have $E_A = E$, and therefore the causal graph of the returned task Π is exactly G. □

In order to create a PDDL task, the SAS$^+$ task is then translated to the STRIPS fragment of PDDL, ignoring the facts that correspond to the last value of the variables representing the at-most-one case. PDDL preconditions are taken from prevails and preconditions, add effects are taken from the effects, and delete effects are taken from the preconditions. Note that if the tasks are translated from STRIPS back to SAS$^+$, there is nothing that enforces that the same mutex groups will be detected. Thus, the causal graph structure is not necessarily preserved by translating to STRIPS and back to SAS$^+$.

Experimental Evaluation

We start by constructing the benchmark set, as described in the previous section. Our benchmark set was generated as follows. For each of the causal graph structures mentioned above, and a value in $[0.1, 0.25, 0.5, 0.75]$ for edge probability (if needed), we create a domain. This results in 27 domains in our case, with 7 causal graph structures that do not consider edge probabilities and 5 causal graph structures that do. For each such domain, we generate 512 instances by uniformly choosing the number of atoms (4 variants), variables (4 variants), goal variables (4 variants), maximum prevail size (2 variants), maximum effect size (2 variants), and the upper bound on the minimum number of atoms per layer (2 variants). Thus, our constructed benchmark set consists of 13824 generated planning tasks. The benchmark set and the code for generating planning tasks according to a given causal structure will be publicly available.

Collection	LM-cut	SYMBA*
bipartite	**128**	48
bidirectional bipartite	**28**	11
chain 0.1	**277**	270
chain 0.5	**285**	267
chain 0.25	270	**272**
chain 0.75	206	**228**
complete	**12**	5
dag 0.1	**129**	63
dag 0.5	**86**	23
dag 0.25	**53**	1
dag 0.75	**67**	9
directed chain	278	**280**
fork	**308**	275
inverted fork	284	**327**
polytree 0.1	**254**	206
polytree 0.5	**285**	244
polytree 0.25	**298**	233
polytree 0.75	**284**	231
random 0.1	**77**	15
random 0.5	**32**	1
random 0.25	**56**	19
random 0.75	**7**	0
star 0.1	**164**	103
star 0.5	**189**	104
star 0.25	**177**	103
star 0.75	**154**	111
tree	**304**	233
Sum (13824)	**4692**	3682

Table 1: Overall coverage per domain, A^* with LM-cut heuristic and SYMBA* planner. Bolded results indicate the better coverage.

Our evaluation of the constructed set aims at understanding whether the set is sufficiently challenging for modern cost-optimal planners. As there is no clear dominance between symbolic search and heuristic search planners, we have decided to choose one planner for each of these approaches. For symbolic search, we chose the state-of-the-art SYMBA* planner (Torralba et al. 2014). For heuristic

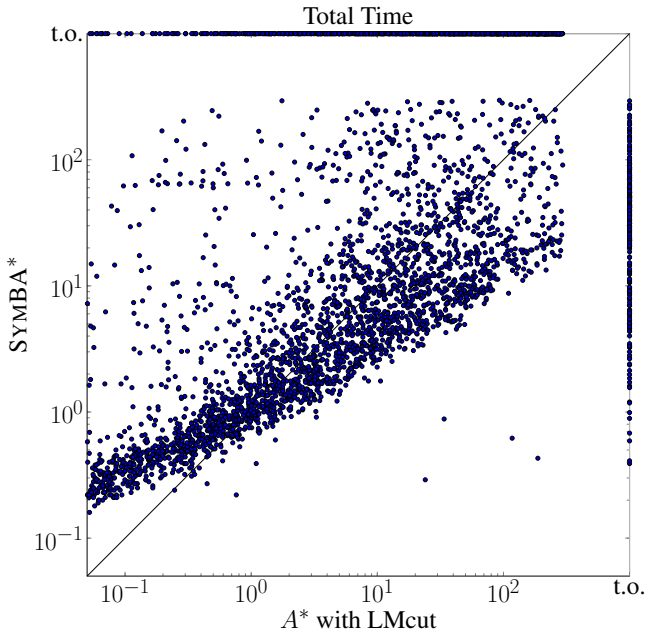

Figure 2: Total time comparison of A^* with LM-cut heuristic to SymBA* planner.

search, we chose A^* with LM-cut heuristic (Helmert and Domshlak 2009), one of the most informative admissible heuristics for classical planning. These two planners are run on the entire constructed benchmark set. As the benchmark set consists of a large number of tasks, we limit the time allocated to each planner to 5 minutes. The experiments were performed on Intel(R) Xeon(R) CPU E7-8837 @2.67GHz machines, with a memory limit of 7.6GB allocated to each run.

Table 1 shows per-domain coverage comparison of the two planners. The coverage captures only the solvable tasks, while nothing in our task generation approach forces the resulting task to be solvable. In addition, LM-cut was able to recognize unsolvable tasks in 137 cases and SymBA* in 507 cases. Out of these, unsolvability is recognizes in the preprocessing step by LM-cut in 24 cases, while by SymBA* in all 507 cases. Note that for each of the techniques, on each of the domains, there still remains a significant number of tasks not solved. Further, there is a significant number of tasks in each domain that were not solved by any of the two techniques. Going beyond aggregated coverage results, Figure 2 shows the total time it takes each planner on each of these tasks. Note that there is no clear advantage to any of the approaches. Looking at the timeouts, there are 1380 cases where SymBA* times outs but LM-cut does not, and 233 cases where LM-cut times outs but SymBA* does not. Clearly, the generated tasks are challenging for both heuristic search and symbolic search based planners.

Related Work

The idea of generating domain models as well as specific planning tasks has been explored in planning community, with a major focus on learning domain models from traces, for classical planning (e.g., (Yang, Wu, and Jiang 2007; Zhuo et al. 2010; Tian, Zhuo, and Kambhampati 2016)) and HTN planning (e.g., (Hogg, Muñoz-Avila, and Kuter 2008; Hogg, Kuter, and Muñoz-Avila 2010; Hogg, Muñoz-Avila, and Kuter 2016)). The work on learning domain models often assumes an existence of a complete model where the plan traces or plan examples are generated from. Some aspects of these domain models are then learned or reconstructed from successful plan traces. Some examples include learning action preconditions (Zhuo et al. 2009), or refine incomplete action descriptions (Zhuo, Nguyen, and Kambhampati 2013).

Probably the most related to our current work is the work on generating problem instances for CSP/SAT problems (e.g., (Achlioptas et al. 2000; Xu et al. 2005)). There are also several online tools/services such as the "Tough SAT Project" or "SATLIB" that generate CNF formulas encoding "difficult" problems (e.g., (Yuen and Bebel 2017; Hoos and Stützle 2000)). Producing hard satisfiable instances has several advantages one of which is to advance the research field in SAT/CSP by providing a suite of problems that can be used for evaluation of solvers. Further, these instances can be polynomial-time reduced to STRIPS in theory, but also in practice (Porco, Machado, and Bonet 2011). The authors provide a tool to translate multiple NP-complete computational problem instances (including SAT, CLIQUE, DirectedHamiltonianPath, etc.) into an NP-Complete fragment of STRIPS that they call STRIPS-1. In that fragment, the actions are either delete-free or can be applied at most once. While the fragment is somewhat limited, the approach can be used for creating additional benchmark sets for planning. Unfortunately, the work has not yet received the attention it deserves, and the instances or the tool are not currently widely used.

It is worth mentioning that our suggested approach to generating random PDDL instances is a somewhat different task than generating random CNF formula, and then translating to STRIPS. Our focus is on being able to control the causal structure of the generated problem, which is not possible with the aforementioned methods.

Summary

In this work, we have presented an approach that allows to generate planning tasks with the causal graph of a specific given structure. Further, we cast these tasks into a STRIPS fragment of PDDL, allowing using as an input to any PDDL planner that supports the STRIPS fragment, as most mature planning systems do. We have generated a collection of 27 domains characterized by the causal graph structure, with 512 tasks in each domain, summing up to 13824 ground PDDL tasks in total. Our experimental evaluation clearly shows that the generated benchmark set is challenging for both the heuristic search based and the symbolic search based planners. In the hope to facilitate further research and enable better comparison of planning tools, we make our tool publicly available to the planning community.

For future work, we intend to explore additional structural restrictions of planning tasks, such as, e.g., k-dependence

(Katz and Domshlak 2007), as well as possibly additional causal graph structures. Further, we would like to explore the possibility of generating lifted PDDL tasks of a meaningful causal structure. For that, we would need to understand how to characterize the ground concept of causal graph structure on a lifted level.

References

Achlioptas, D.; Gomes, C. P.; Kautz, H. A.; and Selman, B. 2000. Generating satisfiable problem instances. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, 256–261.

Bäckström, C., and Nebel, B. 1995. Complexity results for SAS$^+$ planning. *Computational Intelligence* 11(4):625–655.

Boddy, M.; Fox, M.; and Thiébaux, S., eds. 2007. *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling (ICAPS 2007)*. AAAI Press.

Bonet, B.; Palacios, H.; and Geffner, H. 2009. Automatic derivation of memoryless policies and finite-state controllers using classical planners. In Gerevini et al. (2009), 34–41.

Domshlak, C., and Brafman, R. I. 2002. Structure and complexity in planning with unary operators. In Ghallab, M.; Hertzberg, J.; and Traverso, P., eds., *Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling (AIPS 2002)*, 34–43. AAAI Press.

Edelkamp, S. 2001. Planning with pattern databases. In Cesta, A., and Borrajo, D., eds., *Proceedings of the Sixth European Conference on Planning (ECP 2001)*, 84–90. AAAI Press.

Fikes, R. E., and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2:189–208.

Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds. 2009. *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS 2009)*. AAAI Press.

Giménez, O., and Jonsson, A. 2008. The complexity of planning problems with simple causal graphs. *Journal of Artificial Intelligence Research* 31:319–351.

Giménez, O., and Jonsson, A. 2009. Planning over chain causal graphs for variables with domains of size 5 is NP-hard. *Journal of Artificial Intelligence Research* 34:675–706.

Grastien, A., and Scala, E. 2018. Sampling strategies for conformant planning. In de Weerdt, M.; Koenig, S.; Röger, G.; and Spaan, M., eds., *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*, 97–105. AAAI Press.

Haslum, P.; Botea, A.; Helmert, M.; Bonet, B.; and Koenig, S. 2007. Domain-independent construction of pattern database heuristics for cost-optimal planning. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI 2007)*, 1007–1012. AAAI Press.

Helmert, M., and Domshlak, C. 2009. Landmarks, critical paths and abstractions: What's the difference anyway? In Gerevini et al. (2009), 162–169.

Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible abstraction heuristics for optimal sequential planning. In Boddy et al. (2007), 176–183.

Helmert, M. 2004. A planning heuristic based on causal graph analysis. In Zilberstein, S.; Koehler, J.; and Koenig, S., eds., *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, 161–170. AAAI Press.

Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.

Hogg, C.; Kuter, U.; and Muñoz-Avila, H. 2010. Learning methods to generate good plans: Integrating HTN learning and reinforcement learning. In Fox, M., and Poole, D., eds., *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2010)*. AAAI Press.

Hogg, C.; Muñoz-Avila, H.; and Kuter, U. 2008. HTN-MAKER: learning htns with minimal additional knowledge engineering required. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI 2008)*, 950–956. AAAI Press.

Hogg, C.; Muñoz-Avila, H.; and Kuter, U. 2016. Learning hierarchical task models from input traces. *Computational Intelligence* 32(1):3–48.

Hoos, H. H., and Stützle, T. 2000. SATLIB: An online resource for research on SAT. In *SAT*, 283–292. IOS Press.

Katz, M., and Domshlak, C. 2007. Structural patterns of tractable sequentially-optimal planning. In Boddy et al. (2007), 200–207.

Katz, M., and Domshlak, C. 2008. Structural patterns heuristics via fork decomposition. In Rintanen, J.; Nebel, B.; Beck, J. C.; and Hansen, E., eds., *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS 2008)*, 182–189. AAAI Press.

Katz, M., and Domshlak, C. 2010. Implicit abstraction heuristics. *Journal of Artificial Intelligence Research* 39:51–126.

Katz, M., and Keyder, E. 2012. Structural patterns beyond forks: Extending the complexity boundaries of classical planning. In Hoffmann, J., and Selman, B., eds., *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, 1779–1785. AAAI Press.

Katz, M.; Sohrabi, S.; Samulowitz, H.; and Sievers, S. 2018. Delfi: Online planner selection for cost-optimal planning. In *Ninth International Planning Competition (IPC-9): planner abstracts*, 57–64.

Ma, T.; Ferber, P.; Huo, S.; Chen, J.; and Katz, M. 2020. Online planner selection with graph neural networks and adaptive scheduling. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*. AAAI Press.

McDermott, D. 2000. The 1998 AI Planning Systems competition. *AI Magazine* 21(2):35–55.

Palacios, H., and Geffner, H. 2009. Compiling uncertainty away in conformant planning problems with bounded width. *Journal of Artificial Intelligence Research* 35:623–675.

Porco, A.; Machado, A.; and Bonet, B. 2011. Automatic polytime reductions of NP problems into a fragment of STRIPS. In Bacchus, F.; Domshlak, C.; Edelkamp, S.; and Helmert, M., eds., *Proceedings of the Twenty-First International Conference on Automated Planning and Scheduling (ICAPS 2011)*, 178–185. AAAI Press.

Sievers, S.; Katz, M.; Sohrabi, S.; Samulowitz, H.; and Ferber, P. 2019. Deep learning for cost-optimal planning: Task-dependent planner selection. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*, 7715–7723. AAAI Press.

Sohrabi, S.; Riabov, A. V.; Katz, M.; and Udrea, O. 2018. An AI planning solution to scenario generation for enterprise risk management. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*, 160–167. AAAI Press.

Tian, X.; Zhuo, H. H.; and Kambhampati, S. 2016. Discovering underlying plans based on distributed representations of actions. In Jonker, C. M.; Marcella, S.; Thangarajah, J.; and Tuyls, K., eds., *Proceedings of the 2016 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2016)*, 1135–1143. ACM.

Torralba, Á.; Alcázar, V.; Borrajo, D.; Kissmann, P.; and Edelkamp, S. 2014. SymBA*: A symbolic bidirectional A* planner. In *Eighth International Planning Competition (IPC-8): planner abstracts*, 105–109.

Xu, K.; Boussemart, F.; Hemery, F.; and Lecoutre, C. 2005. A simple model to generate hard satisfiable instances. In Biundo, S.; Myers, K.; and Rajan, K., eds., *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS 2005)*, 337–342. AAAI Press.

Yang, Q.; Wu, K.; and Jiang, Y. 2007. Learning action models from plan examples using weighted max-sat. *Artificial Intelligence* 171(2-3).

Yuen, H., and Bebel, J. 2017. Tough SAT project.

Zhuo, H. H.; Hu, D. H.; Hogg, C.; Yang, Q.; and Muñoz-Avila, H. 2009. Learning HTN method preconditions and action models from partial observations. In Boutilier, C., ed., *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, 1804–1810. AAAI Press.

Zhuo, H. H.; Yang, Q.; Hu, D. H.; and Li, L. 2010. Learning complex action models with quantifiers and logical implications. *Artificial Intelligence* 174(18).

Zhuo, H. H.; Nguyen, T.; and Kambhampati, S. 2013. Refining incomplete planning domain models through plan traces. In Rossi, F., ed., *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*. AAAI Press.