Automatic Configuration of Benchmark Sets for Classical Planning

Abstract

The benchmarks from previous International Planning Competitions are commonly used to evaluate new planning algorithms. Since this set has grown organically over the years, it has several flaws: it contains duplicate tasks, trivially solvable domains, unsolvable tasks and tasks with modelling errors. Also, different domain sizes complicate the aggregation of results. Most importantly, however, the range of task difficulty is very small in many domains. We propose an automated method for creating benchmarks that solves these issues. To find a good scaling in difficulty, we automatically configure the parameters of benchmarks domains. We show that the resulting benchmark set improves empirical comparisons by allowing to differentiate between planners more easily.

Introduction

The area of planning is concerned with the development of general solvers that find solutions to problems of any kind. This makes the evaluation of planners an essential part of planning research. The International Planning Competition (IPC) has set some evaluation standards and triggered the development of tools that compare planners in terms of different metrics (Linares López *et al.* 2013; Seipp *et al.* 2017; Vallati *et al.* 2018). The most popular metric is coverage, i.e., the number of solved instances within certain time and memory limits. Typically, there are two main goals for the evaluation: (1) analyze the impact of the novel algorithms by comparing their performance against a baseline, and (2) compare the performance against the state of the art to evaluate the progress in the area. Evaluating planners on different benchmark sets may produce different results, leading to different conclusions about the evaluation. Therefore, having a standardized benchmark set is important to increase the comparability of results across different papers, and to avoid the use of benchmarks tailored for the proposed technique.

We focus on classical planning where the current standard benchmark set has grown across the nine editions of the IPC so far, from 1998 to 2018 (Hoffmann and Edelkamp 2005; Linares López *et al.* 2015). Numerous researchers have contributed to this set by carefully considering how to design new domains (Hoffmann *et al.* 2006), so it features a diverse set of domains that pose interesting challenges for planning algorithms. However, there are several issues with this benchmark set (Moraru and Edelkamp 2019). For example, it uses a different number of instances per domain, which reduces the value of statistics aggregated overy different domains. Moreover, instances in the current benchmark set were scaled to be useful for the evaluation of planners at the given IPC. Some of the domains are trivially solved by modern planners, making it impossible to show any coverage advantages over a baseline. On the other hand, early IPC editions did not have a specialized track for optimal planning, and some of their instances are too hard even for state-of-the-art optimal planners.

This paper deals with the problem of how to generate instances of a domain to evaluate planning algorithms. Our goal is to improve the empirical evaluation of future planning papers by (1) providing an algorithm for automatically constructing interesting benchmark sets and by (2) using this algorithm to construct a new benchmark set where differences in performance are better reflected in coverage than under the current standard. We aim to generate a set of instances that range from very easy (solved by most planners) to very hard (out of reach for current state-of-the-art planners) allowing future approaches to show benefits with respect to the harder instances. This definition necessarily depends on the algorithms being evaluated.

We identify which properties are desirable for a benchmark set and propose an automatic method that generates a set of instances, given an *instance generator*, a *baseline* planner that represents the expected minimum performance of any planner, and a set of *state-of-the-art* planners. The instance sets generated by our method fulfill the desirable properties by design. To avoid overfitting to the sets of planners used and not introduce a bias in our benchmark set, our method does not select a set of instances directly, but rather performs a search on the space of possible parameters for the generator to obtain a set of instances of adequate difficulty. We use our tool to design two separated sets of benchmarks, for optimal and satisficing planning, and show their advantages over the current IPC standard.

Background

Informally, a classical *planning task* is defined by an initial state, a set of actions and a goal description. Given a planning task, a *planner* finds a *plan*, that is, a sequence of actions that can be applied in the initial state to achieve the

goal. A plan is optimal if it minimizes the summed-up cost of the actions among all plans. If the planner is guaranteed to find an optimal solution, it is an optimal planner, otherwise it is a satisficing planner.

Since its inception in 1998, the International Planning Competition (IPC) has set the standards for the evaluation of planners such as the planner input language PDDL (McDermott *et al.* 1998). The IPC also introduced numerous planning tasks from different problem settings, called *domains*.

A planning task is typically divided into a domain and an instance file. The domain file defines the types of objects, their properties, and the action schemas. Each instance file can have a different number of objects, initial state and goals. Most domains have an instance generator, a program that, given certain parameters and a random seed, will generate a new instance of the domain. Even though many instance generators are available,[1] most planning papers use the benchmarks introduced for the IPCs, since a standardize benchmark improves the reproducibility of planning research.

As an example, consider the Nomystery domain, where a truck must deliver a set of packages to certain locations. To do that, there is a limited amount of fuel that is consumed by drive actions. Instances differ on the amount of fuel available, the number of locations and their connections, the number of packages, and their initial and final location. The instance generator for Nomystery accepts several parameters that allow the benchmark designer to control the difficulty of the generated instances: the number of locations, the number of packages, the number of edges between locations, the maximum fuel consumption between two locations, and the constrainedness $C \geq 1$, so that the amount of fuel in the initial state is set to C times the minimum fuel consumption required to solve the instance.

Benchmark Design Principles

The purpose of a benchmark set is to evaluate planners and compare their performance on a diverse class of problems. In this section, we analyze which properties are desirable of an instance set so that the evaluation of a set of planners is as informative as possible. For aggregated statistics to be meaningful, not only should all domains have the same number of instances, but their difficulty should also scale similarly. A successful scaling must meet three conditions: (1) have easy instances that are solved by all planners, (2) have hard instances that are not solved by any current planner, and (3) the instance difficulty should grow smoothly.

Condition (1) is necessary for experiments to be informative at all: if some planners do not solve any instance, no conclusions can be obtained about their relative performance. Condition (2) is necessary for new algorithms to show that they can deal with instances that previous planners could not. Condition (3) is necessary for differences between the planners' performance to be reflected in coverage. To show why, consider an idealized setting where a baseline planner A, whose runtime scales exponentially ($t(A, x) = x^C$ for some constant C), is compared to an improved planner version B which is always faster than A by a factor of $K \geq 0$, i.e., $t(B, x) \leq \frac{t(A,x)}{K}$. Given these assumptions, there is a guaranteed difference in coverage if and only if (1) some instances are solved by B, (2) not all instances are solved by A, and (3) $K \geq C$. Otherwise, there may be cases where both planners solve the same number of instances, and the difference in performance by a factor of K is missed by the coverage analysis. For example, if $K = 2$ and $C = 3$, then (3) does not hold. So for any time limit (e.g., 300 seconds), if the runtime of the last instance solved by A is close enough to the time limit (e.g., 250 seconds), the next instance cannot be solved by B below the time limit (e.g., $\frac{250 \cdot 3}{2} > 300$).

Real distributions of planner runtimes over sets of instances differ from this idealized example in that they usually involve constant factors, the runtime scaling of different planners may be completely different, and even for a single planner it may be impossible to obtain instances that scale according to the desired runtimes in some domains. But ideally, all domains should consist of a collection of instances of increasing difficulty, ranging from very easy to very hard for current planners. Therefore, we aim for a collection where the easiest instance is quickly solved by most planners; all domains have instances that are not solved by current state-of-the-art planners; and difficulty scales by approximately a factor of 1.5–2 between consecutive instances.

Figure 1 exemplifies some of the issues of the IPC benchmark set. The left plot shows the results of three standard optimal planners on the Barman domain from IPC 2011 and our new benchmark set (New'14). In the IPC instances the difficulty does not grow smoothly. Instead, for each group of four instances the difficulty increases visibly and the runtime of all planners increases by about one order of magnitude. This is undesirable since we cannot observe differences in performance for some planners by inspecting their coverage. For example, Blind is six times faster than LM-cut, but both solve four instances in 30 minutes. Only if we lower the time limit to 30–205 seconds, Blind solves more tasks than LM-cut. In contrast, the difficulty on the new benchmark set grows more smoothly, there are instances of more varied difficulty for all planners, and fewer jumps in their runtime. Accordingly, for almost all time limits between 1 and 1800 seconds, Blind has a higher coverage than LM-Cut.

The right part of Figure 1 shows an example for satisficing planning using the Zenotravel domain introduced in IPC'00. The IPC'00 instances do not allow to identify which of the three planners is preferable in this domain. All three planners solve all instances below 20 seconds and the runtime difference could very well be due to a constant factor. With this set of Zenotravel instances it is very hard to argue that a new algorithmic idea is superior to LAMA or even to GBFS-FF. The reason is not methodological; the selected instances were adequate to evaluate existing planners in the early 2000's. However, the state of the art in planning has advanced in the last two decades and those instances are not good enough anymore to evaluate modern

[1] https://bitbucket.org/planning-researchers/pddl-generators/

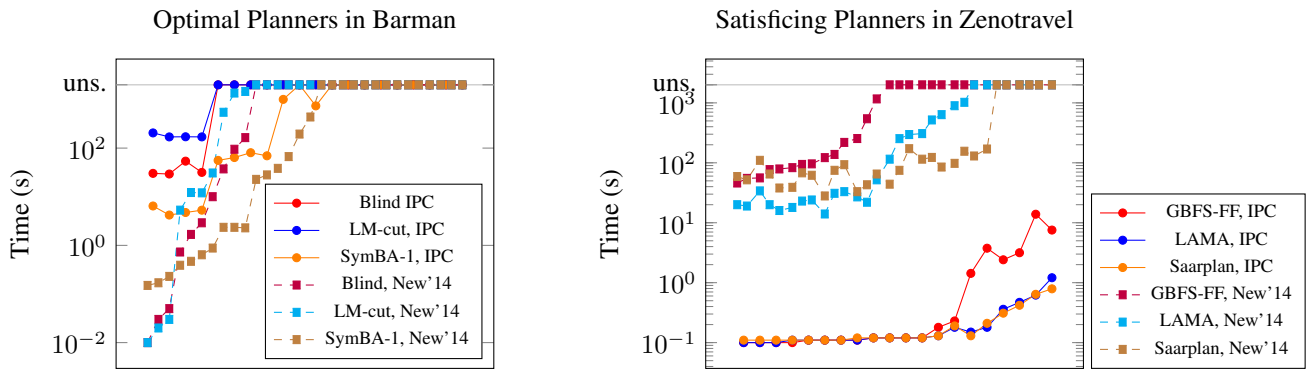

Figure 1: Runtime comparison of optimal and satisficing planners on two domains from the IPC and the New'14 benchmark set.

planners. In constrast, the results with our new benchmark set, which used planners from 2014 and earlier (GBFS-FF and LAMA, but not Saarplan) for selecting the instances, show that Zenotravel can still be an interesting domain to evaluate state-of-the-art planners. With the new set, we can clearly observe both in terms of coverage and runtime that Saarplan is slightly superior to LAMA (especially for harder instances) and both are much faster than GBFS-FF in this domain. Even though the optimization did not take Saarplan into account, the smooth runtime scaling for LAMA makes it easier to observe such a difference.

Given this definition of an ideal benchmark set, the instance selection necessarily depends on the planning algorithms being evaluated. However, one must be careful of not "overfitting" the benchmark set to match the set of selected planners, so that the difficulty scales well over the selected planners but not for future planning algorithms. To avoid overfitting, we impose two restrictions on the benchmark optimization process. On the one hand, the optimization process cannot consider individual instances, so the benchmark set will be generated with a different random seed than the one used during training. On the other hand, we do not use all planners available for the optimization. In each domain, we require a *baseline* planner that represents the expected minimum performance of any planner to ensure that some instances are solved by all planners. Also, to ensure that some instances remain unsolved, we estimate the performance of *state-of-the-art* planners by taking the minimum runtime of any of the considered planners on each instance. This makes our instance selection as objective as possible since it does not depend on the concrete set of planners available to the benchmark designer, as long as the best planner for the given domain is considered.

Optimization of Planning Benchmarks

We consider domains with an instance generator that has several parameters to control the hardness of the generated instances. We distinguish between two types of parameters. *Linear parameters* can be assigned arbitrary non-negative numeric values, where larger values usually result in harder problems. They are typically used to specify the number of objects of a given type. Each generator should have at least one linear parameter that helps to control the difficulty of the generated instances. In contrast, *enumerated parameters* have a finite set of values, and we do not make any assumption over their impact on instance hardness. All other parameters are fixed to a predefined constant value.

We define the instances of a domain as a set of *sequences* of instances. A sequence consists of a list of planning tasks $\Pi_1, \Pi_2, \ldots$ of increasing difficulty. To ensure that difficulty increases, all instances in the sequence have a fixed value for all enumerated parameters, whereas the value of linear parameters increases linearly across the sequence. We specify this via the base value b and slope m that the linear parameter takes on Π_1. For example, suppose that a domain has two linear parameters that define the number of packages ($b = 2, m = 1$), and trucks ($b = 1, m = 0.5$). Then, the sequence will generate instances with the following numbers of packages and trucks: $(2, 1), (3, 1), (4, 2), (5, 2), (6, 3)$, etc.

Considering sequences of instances allows us to make informative decisions about the parameters that should be used to generate instances unsolvable by current planners. A limitation of this approach is that not all combinations of parameters are possible. For example, a single sequence cannot contain both $(3, 2)$ and $(2, 3)$ because that would require to decrease one of the parameters, which is not allowed by our linear scaling. In most cases, this is not a problem because multiple sequences of instances can be used for a single domain. A notable exception are parameters that define the width and height of a grid, because they have a strong interaction, i.e., the number of tiles is the product of both parameters. In that case, we consider them a single parameter so that the number of tiles in the grid scales linearly.

In order to use our benchmark configuration tool, the benchmark designer must specify how to call the instance generator and provide lower and upper bounds for all linear parameters, as shown in the snippet from Figure 2. This allows us to specify preferences on which parameters to scale (e.g. by setting the slope m on the number of locations to be in the $[0.1, 1]$ interval, we indicate our preference by increasing difficulty by scaling the number of packages). Also, any constraint imposed by the generator on the value of its pa-

```
generator_command="nomystery -l {locations} -p {packages}
    -n {edgefactor} -m {edgeweight} -c {constrainedness}
    -s {seed} -e 0 ",
  [LinearAttr("locations", lower_b=3, upper_b=5,
                           lower_m=0.1, upper_m=1),
   LinearAttr("packages", lower_b=2, upper_b=10),
   ConstantAttr("edgefactor", "1.5"),
   ConstantAttr("edgeweight", "25"),
   EnumAttr("constrainedness", [1.1, 1.5, 2.0])]
```

Figure 2: Example of a domain specification with the generator command and the specification of the corresponding parameters.

rameters must be enforced by adding a postprocessing function that updates the value of the parameters passed to the generator. For example, if the number of packages has to be greater than the number of locations, instead of directly selecting the amount of packages, our linear scaling will consider the number of locations and the number of additional packages. All these adjustments must be done on a per-domain basis, since they depend on the specific characteristics of each domain and instance generator.

Given this framework, our automatic tool decides which sequences of instances are suitable for each domain. This is done in two phases: the first phase designs a set of candidate sequences (Sequence Optimization), and the second phase performs a final selection that adheres to our design principles as much as possible (Sequence Selection).

Sequence Optimization

The first phase generates sequences of 30 instances by optimizing sequence parameters. To guide the search towards sequences that scale the instance difficulty as smoothly as possible, we compute a penalty score for each sequence and search for the sequence that minimizes this score.

Sequences are evaluated by running the state-of-the-art ($\mathcal{A}$) and baseline ($\mathcal{B}$) planners on their instances, using a time limit of 180 seconds per instance. We ignore instances that are solved under 10 seconds, considering that differences of ± 5 seconds are not meaningful enough. Since the sequences are generated with increasing values of the linear parameters, we assume that the runtimes will always increase, so we can stop our evaluation as soon as one instance is not solved under the time limit. In cases where this does not hold, we enforce it by sorting the runtimes of the instances. Our assumption is that these anomalies stem from using different random seeds for the instance generator and the results will be different with different random seeds.[2] The runtime of a set of planners is the minimum of the runtimes of the individual planners. Let $t(X,1),\ldots,t(X,5)$ be the runtimes of the set of planners X on the first five instances with a runtime above 10 seconds. The penalty score is defined as $\sum_{i\in[2,5]} S(\mathcal{B},i) + S(\mathcal{A},i)$ where $S(X,i) =$

$$\begin{cases} 3 - \dfrac{2t(X,i)}{t(X,i-1)} & \text{if } 1 \leq \dfrac{t(X,i)}{t(X,i-1)} \leq 1.5 \\ 0 & \text{if } 1.5 < \dfrac{t(X,i)}{t(X,i-1)} \leq 2 \\ 1 - \dfrac{2t(X,i-1)}{t(X,i)} & \text{if } 2t(X,i-1) \leq t(X,i) \leq 180 \\ 2 & \text{if } t(X,i) > 180 \end{cases}$$

This penalty is lower for sequences whose runtime scales smoothly, assigning a minimum score of 0 to any sequence where the runtimes of both the baseline and state-of-the-art planners scale exponentially with a factor between 1.5 and 2, e.g., $\langle 10, 15, 23, 35, 52, \ldots \rangle$, or $\langle 10, 20, 40, 80, 160, \ldots \rangle$. If not enough instances are solved in the $[10, 180]$ second interval, the sequence gets a penalty of 2, and otherwise a penalty between 0 and 1 is assigned. To avoid generating sequences where all instances are solved by the state-of-the-art planners, we also add a penalty of 1 for each instance solved by them beyond 20 instances. To guarantee that all valid sequences contain some solvable instances and to speed up the evaluation we require the first three instances to be solved within 10, 60, and 180 seconds, respectively. Otherwise, we discard the sequence, unless all linear parameters are already taking their minimum base value.

The choice of the concrete penalty values is arbitrary. What matters is that sequences that minimize this score adhere more to the design principles introduced in Section than those that do not, thereby guiding the parameter optimization towards good sequences.

Sequence Selection

After performing one or more optimization runs (using different random seeds) as described above, we collect all sequences seen during the optimization process. As the number can be extremely large, we only keep the 100 sequences with the lowest penalty score per value of the enumerated parameters, eliminating sequences that share all evaluated instances (i.e., those solved by the baseline or state of the art planners) with a previously selected sequence. This filter ensures that we select a set of diverse sequences with a good penalty score.

For each sequence, we run the baseline and state-of-the-art planners to obtain the runtimes of all instances solvable in 180 seconds. For the rest of the instances, we estimate the runtime by assuming that runtimes will increase according to the average increasing factor $t(\mathcal{A},i)/t(\mathcal{A},i-1)$ observed on the instances solved between 5 and 180 seconds. This is a very rough estimate but it is accurate enough for the purposes of choosing up to when a sequence should be continued (see below).

We model the problem of selecting a suitable set of subsequences as a Mixed-Integer Programming problem, where constraints directly aim to model the design principles of Section . The decision variables model the start and end points of each sub-sequence of instances. The selection must satisfy the following *hard constraints* that model properties desirable for a good set of instances:

[2]Note that any parameter whose influence on the planner's performance is unpredictable should be considered an enumerated parameter and remain constant for a given sequence.

Training Optimal	all five components of the FDSS 2 portfolio from IPC 2011 (Helmert *et al.* 2011) and Symba-1 from IPC 2014 (Torralba *et al.* 2014)
Training Satisficing	greedy best-first search + FF heuristic (Hoffmann and Nebel 2001), LAMA (Richter and Westphal 2010) and Madagascar (Rintanen 2012)
Evaluation Optimal	five components of Delfi1 portfolio from IPC 2018 using symmetry pruning and partial order reduction (blind search, iPDB, LM-Cut and two M&S variants, see Katz *et al.* 2018) and three vanilla IPC 2018 planners: Complementary2 (Franco *et al.* 2018), DecStar (Gnad *et al.* 2018), Scorpion (Seipp 2018b)
Evaluation Satisficing	eight vanilla IPC 2018 planners: Cerberus (Katz 2018), BFWS-PREF, DUAL-BFWS and POLY-BFWS (Francès *et al.* 2018), DecStar (Gnad *et al.* 2018), OLCFF (Fickert and Hoffmann 2018), Fast Downward Remix (Seipp 2018a) and Saarplan (Fickert *et al.* 2018)

Table 1: Choice of planners for benchmark generation and evaluation.

(H1) The number of selected instances must be exactly 30.

(H2) There must be at least one instance solvable by the baseline under 30 seconds.

(H3) All sequences must start with an instance that is solvable by a state-of-the-art planner and end with an instance whose estimated runtime is higher than 2000 seconds.

(H4) Each parameter configuration must be used (with different random seeds) at most twice, and only once for domains whose generators do not admit a random seed.

The objective is to minimize the summed-up penalty score of all sequences used, plus the penalty incurred for violating any of the following *soft constraints*:

(S1) The number of instances solved by the baseline under 30 seconds must be between 2 and 6 (with a penalty of $2x^2$ where x is the deviation with respect to the constraint).

(S2) The number of instances solved under 180 seconds must be between 8 and 15 (with a penalty of $2x^2$ where x is the deviation with respect to the constraint).

(S3) All sequences must end with an instance whose estimated runtime is between 18 000 and 180 000 (that is, 1–2 orders of magnitude more than the typical time limit of 30 minutes). Larger times t incur a penalty of $100t/180000$ and smaller times incur a penalty of $100(18000/t)$.

(S4) If a parameter configuration is used more than once, there is a penalty of 100.

On domains where all instances in the sequence are solved by state of the art planners under 180 seconds (because the domain is solvable in polynomial time and it is impossible to fulfill our criteria with state-of-the-art planners), we consider the runtimes of the baseline instead of those of the state of the art planners in our constraints described above.

Constraints (H2), (S1) and (S2) ensure that the instance set contains some easy instances, so that any future planning algorithms are expected to solve some instances, allowing researchers to analyze the behaviour of their algorithms in the domain. Constraints (H3) and (S3) ensure that, whenever possible, at least some of the instances are expected to be out of reach for state-of-the-art planners. Together with minimizing the penalty score of the selected sequences, they aim to obtain a smooth scaling, since sequences must interpolate between easy and hard instances and sequences with smoother scaling are preferred. Finally, constraints (H4) and (S4) are needed to avoid duplicate instances and instances that are very similar to each other.

The penalties are set arbitrarily, but they scale quadratically with respect to the deviation because it is better to not fulfill several soft constraints entirely than to completely ignore one of the constraints.

Experiments

We implemented the first phase, i.e., sequence optimization, using the automatic configurator SMAC (Hutter *et al.* 2011). We test our approach by running two completely separated optimizations for optimal and satisificing planners. As baseline planners, we use blind search for optimal planning and greedy best-first search with the FF heuristic (Hoffmann and Nebel 2001) for satisficing planning, both implemented in Fast Downward (Helmert 2006a).

Both for optimal and satisficing planning, we generate two separate benchmark sets, New'14 and New'20 that differ on the set of state-of-the-art planners available for training. The New'14 set consists of a heterogeneous set of planning algorithms from IPCs 2011 and 2014, listed under "Training Optimal/Satisficing" in Table 1. Thew New'20 set is trained using the same planners plus the ones used in the evaluation ("Evaluation Optimal/Satisificing" in Table 1). Therefore, for New'14, the training and evaluation sets are disjoint.

Since we limit each planner run during the optimization to 3 minutes, we adapt the planners by breaking portfolios into components and by adapting time limits for preprocessing for the training phase. For optimization of each domain, we hand-pick 1–2 planners that perform best in that domain, and we run SMAC 10 times using different random seeds. Each run is limited to 10 hours. After the first phase finishes, we consider all sequences encountered during optimization for the second phase, i.e., sequence selection. We solve the MIP for sequence selection using CPLEX 12.8, which finishes in under 30 seconds for each domain.

We evaluate the new benchmark sets using the aforementioned planners, limiting each run to 30 minutes and 3.5 GiB. In Table 2 we compare the IPC benchmarks to the new benchmark sets for optimal and satisficing planning. The first column in both tables (#IPC) confirms the problematically large differences in domain sizes, a problem our new benchmark set does not share since it has 30 instances for each domain.

		coverage range			comparisons		
Optimal	#IPC	IPC	'14	'20	IPC	'14	'20
barman	34	4–11	10–14	8–11	13	**21**	19
blocksworld	35	18–30	5–12	5–12	18	**25**	24
childsnack	20	0–6	6–15	8–20	13	**18**	**18**
depot	22	5–14	10–21	4–18	24	**27**	25
driverlog	20	7–15	10–30	7–27	22	**27**	24
elevators	50	28–44	8–14	6–16	**26**	24	22
floortile	40	16–34	7–19	7–17	21	22	**24**
gripper	20	8–20	11–30	11–30	**7**	**7**	**7**
hiking	20	13–18	6–18	8–15	23	**24**	**24**
miconic-strips	150	56–143	3–30	3–27	27	**28**	26
nomystery	20	8–20	7–26	4–23	17	**27**	**27**
parking	40	0–15	8–16	7–14	**28**	27	25
rovers	40	6–13	9–30	10–27	**25**	21	7
satellite	36	7–14	24–30	11–29	**21**	**21**	18
scanalyzer	50	21–33	11–19	9–19	**26**	24	**26**
snake	20	7–14	6–11	4–16	22	19	**24**
storage	30	15–18	10–25	9–20	21	**27**	**27**
tpp	30	7–20	5–30	3–30	24	**27**	**27**
transport	70	24–35	8–25	12–23	21	13	**26**
visitall	40	12–30	6–16	10–20	**26**	25	25
woodworking	50	38–50	14–21	8–18	21	**26**	24
zenotravel	20	8–13	4–22	6–30	23	21	**26**

		coverage range			comparisons		
Satisficing	#IPC	IPC	'14	'20	IPC	'14	'20
barman	40	39–40	3–22	11–30	7	22	**27**
blocksworld	35	35–35	6–22	7–28	0	24	**26**
childsnack	20	1–20	17–30	2–24	**27**	24	25
depot	22	22–22	14–29	11–30	0	**26**	**26**
driverlog	20	20–20	30–30	10–20	0	0	**25**
elevators	50	50–50	26–30	13–30	0	7	**24**
floortile	40	4–40	1–14	1–20	17	22	**26**
gripper	20	20–20	29–30	26–30	0	**7**	**7**
hiking	20	10–20	0–30	3–20	24	**27**	24
miconic-strips	150	150–150	30–30	30–30	**0**	**0**	**0**
nomystery	20	12–20	13–30	0–30	23	22	**25**
parking	40	36–40	19–24	10–12	7	**24**	13
rovers	40	38–40	11–30	11–21	7	**26**	25
satellite	36	26–36	15–30	7–17	23	23	**24**
scanalyzer	50	48–50	9–11	10–16	12	17	**21**
snake	20	3–17	9–30	1–11	**27**	22	**27**
storage	30	21–30	8–21	6–29	26	**27**	**27**
tpp	30	29–30	6–26	8–17	12	**27**	26
transport	70	67–70	13–28	17–26	7	**26**	**26**
visitall	40	36–40	19–30	23–30	7	23	**25**
woodworking	50	28–50	4–30	2–30	13	26	**27**
zenotravel	20	20–20	1–25	3–30	0	**27**	23

Table 2: Comparison of the IPC and new benchmark sets for optimal and satisficing planning. The #IPC column shows the number of tasks per domain in the IPC benchmark set, which is always 30 for the new sets. The coverage range shows the minimum and maximum coverage of any planner. In the "comparisons" columns we list how many pairs of planners yield different coverage scores for each benchmark set.

We compare benchmark sets according to two metrics: the range of coverage scores per domain, which allows us to see how many instances are solved by all planners and how many remain unsolved by any of the planners; and the number of comparisons, which quantifies how many differences in the performance of planners are reflected by the coverage score.

In optimal planning the difference between the benchmark sets is rather subtle because difficulty typically scales very fast with increasing instance size. Therefore, the IPC set has some interesting instances in all domains. Also, it can be very hard to generate instance sequences whose difficulty scales smoothly, since often increasing one of the parameters of a generator by a unit has a big impact on the runtime of planners.

The results are more pronounced for satisficing planners, where the IPC set scales very poorly for some domains. Only in two domains the IPC set is slightly superior in terms of comparisons detected by the coverage score, and they are domains manually scaled by the IPC organizers recently: Childsnack from IPC'14 and Snake from IPC'18. In contrast, with the new benchmark sets, we observe differences in performance in domains like Blocksworld, Depot or Zenotravel, where all planners solve all instances in the IPC set. Overall, New'14 uncovers more differences in coverage between pairs of planners than the IPC set in 16 domains, while the opposite is the case in only 3 domains.

The comparison between New'14 and New'20 reveals that our technique is not very sensitive to the set of state-of-the-art planners. The reason is that the state of the art has not advanced enough in the last six years to make a set of instances trained with our method in 2014 outdated.

Discussion

Our paper deals with the problem of generating instances that are adequate to evaluate planning algorithms. The goal is to select instances that scale well, so that differences in algorithm performance are reflected in the number of problems solved within a certain time limit. It must be remarked that no benchmark set can replace a careful analysis of the results. Aggregating results from different domains without further analysis may be misleading and an empirical analysis only based on total coverage should be discouraged. Nevertheless, coverage is a useful metric to summarize experimental data and it is used by most planning papers. As shown by our experiments, the coverage metric is more meaningful for the benchmark sets generated with our approach than with the previous standard. Our main result is a new benchmark set, as well as a set of generators and tools that can be used

in the future to automatically generate new instances.[3]

In other communities like SAT, there has been a lot of research on how to construct random instances (Selman *et al.* 1996; Achlioptas *et al.* 2000; Giráldez-Cru and Levy 2015; Xu *et al.* 2005) around the phase transition (Cheeseman *et al.* 1991). Our approach is orthogonal to any approach that can generate new instances, e.g., around the phase transition of planning problems (Rieffel *et al.* 2014), or with suitable initial states and goals for Sokoban (Bento *et al.* 2019). Also, it could be adapted to generate benchmark sets with different characteristics, e.g., with smaller instances that can be solved in a few seconds (Ruml 2010). Future work could also consider the relation between different domains theoretically (Helmert 2003; 2006b) or empirically (Cenamor and Pozanco 2019).

References

Dimitris Achlioptas, Carla P. Gomes, Henry A. Kautz, and Bart Selman. Generating satisfiable problem instances. In *Proc. AAAI 2000*, pages 256–261, 2000.

Dâmaris S. Bento, André Grahl Pereira, and Levi H. S. Lelis. Procedural generation of initial states of sokoban. In *Proc. IJCAI 2019*, pages 4651–4657, 2019.

Isabel Cenamor and Alberto Pozanco. Insights from the 2018 IPC benchmarks. In *ICAPS Workshop on the IPC (WIPC)*, pages 8–14, 2019.

Peter Cheeseman, Bob Kanefsky, and William M. Taylor. Where the really hard problems are. In *Proc. IJCAI 1991*, pages 331–337, 1991.

Maximilian Fickert and Jörg Hoffmann. OLCFF: Online-learning h^{CFF}. In *IPC-9 planner abstracts*, pages 17–19, 2018.

Maximilian Fickert, Daniel Gnad, Patrick Speicher, and Jörg Hoffmann. SaarPlan: Combining saarlands greatest planning techniques. In *IPC-9 planner abstracts*, pages 11–16, 2018.

Guillem Francès, Hector Geffner, Nir Lipovetzky, and Miquel Ramiréz. Best-first width search in the IPC 2018: Complete, simulated, and polynomial variants. In *IPC-9 planner abstracts*, pages 23–27, 2018.

Santiago Franco, Levi H. S. Lelis, and Mike Barley. The Complementary2 planner in the IPC 2018. In *IPC-9 planner abstracts*, pages 32–36, 2018.

Jesús Giráldez-Cru and Jordi Levy. A modularity-based random SAT instances generator. In *Proc. IJCAI 2015*, pages 1952–1958, 2015.

Daniel Gnad, Alexander Shleyfman, and Jörg Hoffmann. DecStar STAR-topology DECoupled search at its best. In *IPC-9 planner abstracts*, pages 42–46, 2018.

Malte Helmert, Gabriele Röger, Jendrik Seipp, Erez Karpas, Jörg Hoffmann, Emil Keyder, Raz Nissim, Silvia Richter, and Matthias Westphal. Fast Downward Stone Soup. In *IPC 2011 planner abstracts*, pages 38–45, 2011.

Malte Helmert. Complexity results for standard benchmark domains in planning. *AIJ*, 143(2):219–262, 2003.

Malte Helmert. The Fast Downward planning system. *JAIR*, 26:191–246, 2006.

Malte Helmert. New complexity results for classical planning benchmarks. In *Proc. ICAPS 2006*, pages 52–61, 2006.

[3]Links to new benchmarks and tool repository will be added.

Jörg Hoffmann and Stefan Edelkamp. The deterministic part of IPC-4: An overview. *JAIR*, 24:519–579, 2005.

Jörg Hoffmann and Bernhard Nebel. The FF planning system: Fast plan generation through heuristic search. *JAIR*, 14:253–302, 2001.

Jörg Hoffmann, Stefan Edelkamp, Sylvie Thiébaux, Roman Englert, Frederico dos Santos Liporace, and Sebastian Trüg. Engineering benchmarks for planning: the domains used in the deterministic part of IPC-4. *JAIR*, 26:453–541, 2006.

Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *Proc. LION 2011*, pages 507–523, 2011.

Michael Katz, Shirin Sohrabi, Horst Samulowitz, and Silvan Sievers. Delfi: Online planner selection for cost-optimal planning. In *IPC-9 planner abstracts*, pages 57–64, 2018.

Michael Katz. Cerberus: Red-black heuristic for planning tasks with conditional effects meets novelty heuristic and enchanced mutex detection. In *IPC-9 planner abstracts*, pages 47–51, 2018.

Carlos Linares López, Sergio Jiménez Celorrio, and Malte Helmert. Automating the evaluation of planning systems. *AI Communications*, 26(4):331–354, 2013.

Carlos Linares López, Sergio Jiménez Celorrio, and Angel García Olaya. The deterministic part of the seventh international planning competition. *AIJ*, 223:82–119, 2015.

Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. PDDL – The Planning Domain Definition Language – Version 1.2. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, Yale University, 1998.

Ionut Moraru and Stefan Edelkamp. Benchmarks old and new: How to compare domain independence forcost-optimal classical planning? In *ICAPS Workshop on the IPC (WIPC)*, pages 36–39, 2019.

Silvia Richter and Matthias Westphal. The LAMA planner: Guiding cost-based anytime planning with landmarks. *JAIR*, 39:127–177, 2010.

Eleanor G. Rieffel, Davide Venturelli, Minh Do, Itay Hen, and Jeremy Frank. Parametrized families of hard planning problems from phase transitions. In *Proc. AAAI 2014*, pages 2337–2343, 2014.

Jussi Rintanen. Planning as satisfiability: Heuristics. *AIJ*, 193:45–86, 2012.

Wheeler Ruml. The logic of benchmarking: A case against state-of-the-art performance. In *Proc. SoCS 2010*, pages 142–143, 2010.

Jendrik Seipp, Florian Pommerening, Silvan Sievers, and Malte Helmert. Downward Lab. https://doi.org/10.5281/zenodo.790461, 2017.

Jendrik Seipp. Fast Downward Remix. In *IPC-9 planner abstracts*, pages 74–76, 2018.

Jendrik Seipp. Fast Downward Scorpion. In *IPC-9 planner abstracts*, pages 77–79, 2018.

Bart Selman, David G. Mitchell, and Hector J. Levesque. Generating hard satisfiability problems. *AIJ*, 81(1–2):17–29, 1996.

Álvaro Torralba, Vidal Alcázar, Daniel Borrajo, Peter Kissmann, and Stefan Edelkamp. SymBA*: A symbolic bidirectional A* planner. In *IPC-8 planner abstracts*, pages 105–109, 2014.

Mauro Vallati, Lukás Chrpa, and Thomas Leo McCluskey. What you always wanted to know about the deterministic part of the international planning competition (IPC) 2014 (but were too afraid to ask). *Knowledge Engineering Review*, 33, 2018.

Ke Xu, Frederic Boussemart, Fred Hemery, and Christophe Lecoutre. A simple model to generate hard satisfiable instances. In *Proc. IJCAI 2005*, pages 337–342, 2005.