

Scaling up ML-based Black-box Planning with Partial STRIPS Models

Matias Greco¹, Álvaro Torralba², Jorge A. Baier¹, Hector Palacios³

¹ Pontificia Universidad Católica de Chile, Chile

² Aalborg University, Denmark

³ ServiceNow Research, Canada

mogreco@uc.cl, alto@cs.aau.dk, jabaier@ing.puc.cl, hector.palacios@servicenow.com

Abstract

A popular approach for sequential decision-making is to perform simulator-based search guided with Machine Learning (ML) methods like policy learning. On the other hand, model-relaxation heuristics can guide the search effectively if a full declarative model is available. In this work, we consider how a practitioner can improve ML-based black-box planning on settings where a complete symbolic model is not available. We show that specifying an incomplete STRIPS model that describes only part of the problem enables the use of relaxation heuristics. Our findings on several planning domains suggest that this is an effective way to improve ML-based black-box planning beyond collecting more data or tuning ML architectures.

1 Introduction

Multiple AI areas deal with sequential decision-making problems. When a black-box simulator is available, an increasingly popular approach is to use search guided with *learned* policies (Williams 1992) or value/heuristics functions (Silver et al. 2016), either by ML using collected data or by online reinforcement learning (RL) (Sutton and Barto 2018). The main advantage of these approaches is that the simulator can implement an arbitrarily complex transition function, as a symbolic model is not required to guide the search. However, such ML methods sometimes feature high sample complexity or fail to generalize to instances beyond the training data distribution, running out of time or misleading the search. For improving them, practitioners might attempt to collect more data or to tune the ML algorithms.

An alternative is to use a state-of-the-art automated planner, which leverages a full symbolic description of the state space, e.g., expressed in the *Planning Domain Definition Language* (PDDL) (McDermott 2000). This results in a symbolic model (e.g. in STRIPS formalism (Fikes and Nilsson 1971)) that compactly describes the state space, and can be used to guide the search via relaxation-based heuristics (Bonet and Geffner 2001; Hoffmann and Nebel 2001). In fact, recent comparisons have shown that, whenever a declarative description is available, these model-relaxation heuristics often outperform ML-based heuristics even when learning for a fixed planning task (Ferber, Helmert, and Hoffmann 2020).

In this paper, we are interested in settings where plans are obtained by searching using an ML-based method and a full symbolic model of the black-box simulator is not available. Our hypothesis is that such systems can be enhanced by using an incomplete STRIPS model and combining the guidance of the learned heuristics/policies with relaxation-based heuristics (e.g. FF (Hoffmann and Nebel 2001)) over the partial, incomplete, model. We consider situations where a hypothetical practitioner has already trained a heuristic and/or policy to find sequences of actions that reach a goal. Then, to improve performance, the practitioner models part of the domain in STRIPS/PDDL, resulting in an incomplete model that describes only part of the state-space without assuming that the black-box state is fully symbolic. The question we tackle is, can the practitioner improve the performance of the overall system by specifying only part of the model in a symbolic way, while still relying on the ML-guidance for capturing other (e.g., non-symbolic) aspects? And if so, how important it is which parts of the model are specified?

We take a first step towards answering these questions by running several case studies on planning domains for which a full model is available. This allows us to exhaustively analyze the performance under a diverse range of partial models that focus on different aspects of the domain and compare the relative improvement with respect to the best possible scenario: having a full STRIPS model. Moreover, we also considered cases where the data was the most convenient for training the chosen ML-based baseline. Our methodology is similar to ML papers that evaluate methods where the data is obtained from a known generative model.

Our experiments show that partial STRIPS models can indeed improve the performance over ML-based heuristics/policies that show low coverage when used in isolation. This is the case even when both search guidance mechanisms are combined in a simple way, not requiring a heavily engineered algorithm nor any additional training, which could be expensive. While the overall performance improves as models become more accurate, it sometimes may suffice to specify a small part of the domain to achieve a noticeable improvement. Therefore, our findings suggest that specifying partial symbolic models is a good alternative way to improve ML-based black-box planning beyond collecting more data or tuning ML architectures.

2 Background / Planning Representation

A *planning task* is a tuple $\Pi = \langle S, A, f, s_I, S_G \rangle$, where S is a finite set of states, A is a finite set of actions, $f : S \times A \rightarrow S$ is the transition function, $s_I \in S$ is the initial state and $S_G \subseteq S$ is the set of goal states. A plan $\pi = \langle a_0, \dots, a_n \rangle$ is a sequence of actions from s_0 to any state $s_g \in S_G$ such that $f(f(f(s_0, a_0), a_1), \dots, a_n) = s_g$.

A planning task can be specified in multiple ways. Following Katz, Moshkovich, and Karpas (2018), a *black-box planning task* is a tuple $\Pi_B = \langle S, A, s_I, succ, goal \rangle$, where *succ* and *goal* are arbitrary, black-box, functions, which define $f(s, a) = succ(s, a)$ and $S_G = \{s_g \mid goal(s) = \text{True}\}$, respectively. Black-box tasks are suitable for performing lookahead search, by starting at the initial state and repeatedly applying the *succ* function. Such a search can be guided using heuristics. Formally, a heuristic $h : S \rightarrow \mathbb{R}_0^+$ is a function mapping states to numeric values, such as $h(s)$ estimates the distance from s to the goal. The heuristic function can be implemented manually, or, e.g., learned automatically from previous tasks (Ferber, Helmert, and Hoffmann 2020; Shen, Trevizan, and Thiébaux 2020).

On the other hand, declarative approaches specify planning tasks by representing states in terms of a set of variables or fluents (Fikes and Nilsson 1971; Bäckström and Nebel 1995). A STRIPS task is a tuple $\Pi_D = \langle F, A, s_I, G \rangle$, where F is a set of facts, and A is a set of actions. A state $s \subseteq F$ is a set of facts, $s_I \subseteq F$ is the initial state and $G \subseteq F$ is the goal specification. The semantics of Π_D are defined as follows. $S = 2^F$ is the set of all states, which correspond to sets of facts that are true in such a state. $S_G = \{s_g \mid G \subseteq s_g\}$. An action $a \in A$ is a tuple $\langle pre(a), add(a), del(a) \rangle$, where $pre(a) \subseteq F$ is a set of preconditions, and $add(a) \subseteq F$ and $del(a) \subseteq F$ are sets of add and delete effects, respectively. An action a is applicable in a state s if $pre(a) \subseteq s$. The resulting state of applying an applicable action a in a state s is the state $f(s, a) = (s \setminus del(a)) \cup add(a)$.

STRIPS models are usually expressed in PDDL, where the set of facts F is the set of all possible grounding of a fixed set of predicates over a fixed set of objects (McDermott et al. 1998). Having a STRIPS model enables new ways for automatically obtaining heuristic functions via domain-independent relaxations of the problem, such as the *delete-relaxation* FF heuristic (Hoffmann and Nebel 2001).

3 Guiding Black-Box Searches with Partial STRIPS Models

Our general hypothesis is that symbolic models can help to improve over black-box search methods that are guided with ML-based heuristics and/or policies. We assume a setting, where a practitioner starts with an already trained ML-based method and a black-box simulator, similar to the setting commonly used in Deep RL (DRL) (Arulkumaran et al. 2017). Our hypothesis is that, when ML-based black-box does not scale well, the practitioner can specify a partial STRIPS model and select a suitable search algorithm to lead to further scalability. This hypothesis does not imply that, for a given ML-based method and partial STRIPS model, such an improvement will always manifest. The empirical

results in our evaluation show that such improvements are possible in black-box planning. In the rest of the section we answer questions crucial for proving our hypothesis: what ML-based methods we consider, how can partial STRIPS models be defined, and what search algorithms can be used for combining the strengths of the ML-based methods and the heuristics obtained from the partial STRIPS models.

3.1 ML-based methods

RL methods focus on estimating one of two quantities. One is *value estimation*, that is the expected cost/reward of acting from a state. The second one is *action selection* where an action is chosen for a given state. They are related, respectively, to the standard RL algorithms value-iteration and policy-iteration (Sutton and Barto 2018). For instance, DRL algorithms for learning a policy might use ML for classifying which action should be applied in each state. In such cases, the ML model returns an estimate of the likelihood of the actions. If the estimation were perfect, maximizing the likelihood corresponds to an optimal policy for the problem. When we are using a policy, we will assume that we can rank actions according to its likelihood.

We focus on black-box planning, where these are used to guide a best-first search, either selecting the state with best value or the node with highest likelihood. Therefore, in this work we refer to value estimation as *heuristics*, and action selection as *policies*.

3.2 Partial STRIPS Models

We assume that we are provided a black-box task, Π_B , for which we have limited ML-based search guidance. Our motivation is that practitioners would write down or extract a symbolic model that represents part of the dynamic of the environment. Such a model would be intentionally added, so practitioners have incentives to capture the weakness of their ML-based method. To do so, they will provide a STRIPS task Π_D as well as a way of mapping states from Π_B to Π_D .

Definition 1 (Partial Model). Let $\Pi_B = \langle S, A, s_I, succ, goal \rangle$ be a black-box planning task. A partial model of Π_B is a tuple $\langle \Pi_D, \sigma \rangle$ where $\Pi_D = \langle F, A, \sigma(s_I), G \rangle$ is a STRIPS task, and σ is a mapping $S \mapsto 2^F$ from states in Π_B to states in Π_D .

Partial models allow using any heuristic function that has been defined for STRIPS tasks on searches for Π_B . In particular, let h_D be any heuristic for Π_D , we define a heuristic for Π_B as $h_B(s) = h_D(\sigma(s))$.

In principle, there are no requirements on the relation between Π_B and Π_D . For example, the set of actions in Π_D can be completely different as that of Π_B . If the preconditions and effects of actions in Π_B are hard to formalize in STRIPS, one can model alternative actions encoding each original action as a sequence of STRIPS actions. But of course, the properties of h_B greatly depend on the relation between Π_B and Π_D . An interesting case is whenever Π_D models only part of the state, and the actions in Π_D expressing the same transition function as those in Π_B . In that case, Π_D is a projection of Π_B and the heuristic h_B preserves the properties safeness, admissibility and consistency

of the heuristic h_D over Π_D (Culberson and Schaeffer 1998; Edelkamp 2001).

3.3 Algorithms

We consider two search algorithms that can use both methods we have for guiding the search, the ML-based method and the heuristics from the partial STRIPS models:

1. *Multi-heuristic best-first search* (Helmert 2006) is a variation of *Greedy Best-First Search (GBFS)* with two queues, which interleaves the corresponding queue to select for expansion its best node. We call this algorithm *double-queue*.
2. GBFS with tie-breaking, which selects for expansion the best node according to the main heuristic. If more than one node has the best heuristic value, it uses the secondary heuristic for tie-breaking.

For dealing with policies, we use the concept of *discrepancy* (Harvey and Ginsberg 1995), as was defined by Karoui et al. (2007). The method ranks the successors of a node (starting from 0) in ascending order according to their heuristic value and/or the preferences of the policy. Then, the heuristic value of a node is the sum of the value of its parent plus its rank. Therefore, the value of each node is the sum of all ranks from the initial state to the node resulting on a path-dependent heuristic. This method was previously used in the context of bounded suboptimal search with good results (Araneda, Greco, and Baier 2021; Greco, Araneda, and Baier 2022).

4 Methodology

Even though our hypothesis is that symbolic models can help to improve over ML-based methods in the general case where a complete declarative model is unknown, we evaluate it on a controlled environment using multiple classical planning domains (see Section 5) for which a complete STRIPS model is already available. This allows us to analyze the benefits from using partial models in a systematic way, using several models per domain that include the “original” model that perfectly describes the dynamics of the environment.

4.1 Implementation

All algorithms were implemented on the Fast Downward (FD) Planning System (Helmert 2006). To compute heuristics on partial STRIPS models, we implemented them as domain-dependent task transformations within the planner, allowing the computation of arbitrary heuristic functions on the partial model.

In our experiments, we use the FF heuristic (Hoffmann and Nebel 2001), one of the most popular heuristics for satisficing planning due to its ability for guiding the search towards the goal in a broad class of domains (Hoffmann 2005). Note that our hypothesis is that ML-based methods can be improved by computing *some* relaxation-based heuristic over a partial model. Therefore, our conclusions are not specifically tied to this heuristic and other planning heuristics might be more convenient for other domains.

4.2 ML-based baseline

We found it convenient to use HGN (Shen, Trevizan, and Thiébaux 2020), a method that can produce a high-quality heuristic for specific domains and generalizes well to instances of different size. HGN provides an estimate of the distance to the goal as a real, floating-point number. As Fast Downward uses only integers, we take only 3 decimals of precision, by multiplying the result by 10^3 and rounding to the nearest integer. We trained the heuristic using a 5-fold cross-validation with their default hyperparameters, i.e. hidden size = 32, bins=4 and learning rate = 0.001. We denote this heuristic as h^{hgn} .

We also consider using HGN as a “policy”, by considering that the policy recommends to take successors with lowest heuristic value. As mentioned above, we can handle this by using discrepancy, so that the evaluation of a node depends on the difference in heuristic value with other siblings. We denote this heuristic as π^{hgn} .

In order to analyze how partial declarative models can be useful on scenarios where the NN-based heuristics and/or policies are used in an out-of-distribution scenario, we train HGN on different datasets. In all cases, our datasets consist of 100 instances, which are different from the instances used in our evaluation. The datasets do differ on the way instances are generated. The default variant of our heuristics, h^{hgn} and π^{hgn} , are trained on instances generated by the same generator used for obtaining the test set. The training instances are typically smaller than the test instances, as is common practice in learning approaches for planning. Other heuristics, which we will denote, h^{dataset} or π^{dataset} , use a “biased” training set, where we vary the distribution of initial states and/or goals. In Section 5, we describe the different variants chosen for each domain.

4.3 Metrics

In our analysis, we focus on comparing the performance of all algorithms in terms of number of expansions and do not focus on running time. This is desirable for several reasons. First of all, in our experiments the time spent on the successor generation and/or on computing the model-based heuristics such as FF is negligible to that of evaluating NN-based heuristics such as HGN. Therefore, in terms of runtime it is sometimes desirable not to use the NN-based heuristics. However, we are interested in a setting where using a ML-based heuristic or policy is desirable, e.g., because it is able to capture large parts of the model that are subsymbolic and cannot be considered. Furthermore, runtime measures could be influenced by a number of factors. In practice, different NN-based heuristics might have different computational cost, or specific hardware can accelerate their computation. Besides, whenever the black-box simulator is based on NN predictions (e.g. as in model-based reinforcement learning), the overhead of the successor generation will increase significantly, heavily limiting the amount of nodes that can be expanded within a given time limit.

Therefore, we analyze how both sources of information, h^{hgn} and h^{FF} , can be combined in order to reduce the search effort in terms of node expansions. Focusing on the number of expansions enables a meaningful comparison when

we evaluate not using the NN-based heuristics, so we can assess the power of the partial STRIPS model isolated. And all the conclusions obtained regarding the cases where using partial models can improve the performance in terms of node expansions can be easily extrapolated to runtime, as the computational effort of these heuristics is far from being a bottleneck.

5 Benchmark Domains

We consider three domains from the International Planning Competition (McDermott 2000). We selected domains that are non-homogeneous, meaning that they have objects of different types so that it is natural to model only part of the domain in a declarative fashion by specifying only a sub-set of the actions, predicates and/or object types in PDDL.

5.1 Logistics

Our first domain is the IPC-00 version of Logistics. The task consists of delivering packages from their initial location to their destination on a map that consists of locations grouped in different cities. Across cities, packages must be transported by airplanes, which are only allowed to move in between special locations where there is an airport. Within each city, packages are transported via trucks.

Our partial model, called *air*, only considers transportation by airplane. In *air*, objects represent packages, cities, and airplanes, but there are no trucks and/or locations within the cities. The goal is to have each package in the correct city. Mapping states from the original task to the partial model is straightforward, by mapping the current location of each package and/or airplane into the corresponding city. This models a scenario where the exact location of each object is hard to determine in a symbolic manner, but it is possible to determine a broad area where the object is located.

For training the HGN heuristic, the default dataset (*hgn*) contains 100 instances with 2 to 4 packages, 2 to 5 cities, 2 to 4 airplanes and 2 to 5 locations in each city. We also generate a biased dataset, called *hgn-one*, where instances have the same size but all packages are placed in the same city in the initial state and goal, so it is never necessary to transport them by airplane. The resulting heuristic is expected to approximate the cost of moving trucks, but ignore all actions that involve the airplanes.

The test set contains 50 instances with 3 to 4 packages, 5 to 7 cities, 4 to 5 airplanes and 4 to 6 locations in each city.

5.2 Grid

In the IPC-00 grid domain, a robot can move along a grid transporting keys. The robot can hold one key at a time. If a tile is locked, the robot has to carry a key of the corresponding shape to unlock it. There are actions that move the robot, pick up or leave keys at a given location, and unlock cells. The goal is to place certain keys at specific locations. For this domain, we consider these two partial domains:

1. *robot*: We model only the position and movement of the robot, but no information regarding the keys is modelled,

except which keys have already been dropped at their target. In this case, the planning heuristic estimates the distance the robot should traverse to visit all target locations where keys should be dropped, without having any information regarding the position of the keys. This could be the case in scenarios where a model of the environment of the robot is available but the position of the keys is not available in the symbolic model.

2. *keys*: We only model the position of the keys, but we do not model the position and/or movement of the robot. Therefore, each key can be grabbed and/or left at any position at any time. The heuristic guides the search towards states where more keys have been dropped at their destination. This kind of model could be useful in settings where the movement of the robot is complex to describe in a symbolic manner accurately.

For training, we generate 100 instances of the following datasets, included two biased ones:

1. *default (hgn)*: a grid size between 4x4 to 6x6, 1 to 2 locks and 1 to 3 keys.
2. *hgn-gridsize*: similar to *default*, but all instances have only two positions (grid size of 2x1 tiles) and one is locked. In the instances exists different keys. The robot has to move to the other location trying the correct key. In this case, the trained heuristic can learn that locked positions require to be unlocked by using keys.
3. *hgn-onelock*: similar to *default*, but all instances have a single locked position, even though they may exist multiple keys. Trained heuristics with this data might be sensitive respect to the kind of grid where the robot moves.

Finally, the test set contains 50 instances with grid size between 7x6 to 8x6, 2 to 3 locks and 2 to 4 keys.

5.3 Woodworking with Pickup and Delivery

This is a variation of the IPC-11 domain Woodworking. The domain simulates the work in a woodworking workshop where some wooden pieces are prepared by cutting boards and varnishing, polishing or coloring the resulting pieces, using different tools, such as a grinder, saw, glazer or spray. In our variant of this domain (henceforth Woodworking-PD), there is additionally a road map. All tools are in a location (representing the workshop), but the pieces of wood are scattered across multiple locations and they should be picked up, brought to the workshop, and the prepared pieces should be delivered afterwards. Thus, the domain includes trucks to move the pieces of board. The original version of this domain has action costs, but we consider a unit-cost version, as action costs are not supported by the HGN heuristic. In our formulation, we only include one truck. For this domain, we consider multiple partial models:

1. *wood*: This is the original woodworking model, which ignores the pickup and delivery part.
2. *logistics*: This partial model ignores the woodworking part and considers that all objects are finished but in different places, locating each unfinished object at the workshop. The tasks consists only of moving objects from one location to another.

	none	h_A^{FF}	h^{FF}
none	-	0	50
$h^{hgn-one}$	5	17	50
h^{hgn}	44	50	50
$\pi^{hgn-one}$	3	3	50
π^{hgn}	50	50	50

(a) Logistics

	none	h_R^{FF}	h_K^{FF}	h^{FF}
none	-	34	50	50
$h^{hgn-gridsize}$	17	34	50	50
$h^{hgn-onelock}$	49	49	50	50
h^{hgn}	50	50	50	50
$\pi^{hgn-gridsize}$	33	49	50	50
$\pi^{hgn-onelock}$	50	50	50	50
π^{hgn}	50	50	50	50

(b) Grid

	none	h_L^{FF}	h_W^{FF}	h^{FF}
none	-	2	19	50
$h^{hgn-move}$	14	24	24	50
h^{hgn}	16	11	36	50
$h^{hgn-oneloc}$	26	33	40	50
π^{hgn}	10	11	14	49
$\pi^{hgn-oneloc}$	12	24	10	50
$\pi^{hgn-move}$	17	23	18	49

(c) Woodworking-PD

Table 1: Coverage of GBFS combining variations of h^{hgn} and h^{FF} with the double-queue algorithm. The top segment is coverage using only STRIPS models. The left segment is coverage using only HGN. For comparison, the right segment is coverage using a complete STRIPS model. The middle segments are using HGN as a heuristic. The bottom ones are coverage using HGN as a policy. Variations of h^{hgn} and π^{hgn} indicate the data used for training. Variations of h^{FF} indicate the partial STRIPS model. Cells in bold indicate those cases where the combination outperforms both FF and HGN in isolation. We highlight in red those cases where using a partial STRIPS model is detrimental.

For training, we generate 100 instances of the following datasets, which included two biased ones that consider just the woodworking part or the pickup and delivery part.

1. *default (hgn)*: the instances have 1 to 2 pieces, 2 to 6 locations, 3 to 5 machines and a wood factor (amount of wood with respect to the minimum needed) of 1.4 to 1.0.
2. *hgn-oneloc*: similar to *default*, but with only one location, and all pieces and machines are in there. For that reason, the goal can be achieved without moving the truck, similar to the original woodworking domain.
3. *hgn-move*: similar to *default*, but for achieving the goal, it is necessary to move the corresponding pieces to a location and not consider processing the pieces (such as paint, varnish, polish, etc). Thus, the goal can be achieved by just moving the truck.

Finally, for testing, we generated a dataset with 50 instances with 2 to 3 parts, 3 to 8 locations, and a wood factor of 1.4 to 1.0.

6 Evaluation

We report experiments on how partial STRIPS models improve over ML baselines. All experiments were run with a time limit of 120 minutes, a memory limit of 9GB, and were executed on a cluster with Intel Xeon Gold 6126 CPU with hyper-threading enabled and around 30 tasks in parallel using Downward lab (Seipp et al. 2017). All our code, benchmarks and data are publicly available (Greco et al. 2022b).

For the evaluation, we use 50 instances in each domain, all different from those used for training the heuristics. In all our experiments, we limit the number of expansions by all configurations to 10 000, the number of expansions that all configurations could complete under our time limit.

6.1 Guiding Search with Partial STRIPS Models

Table 1 reports the coverage when using different kinds of HGN heuristics (trained with the original dataset or with bi-

ased ones), and FF heuristics (using either the full or partial STRIPS models).

The top segment of each table is the coverage when using only the planning heuristic over partial STRIPS models. As a reference for comparison, we include full STRIPS model combined with the FF heuristics (h^{FF}). Without having access to any learned heuristic, it solves all instances under 10 000 expansions. That offers an upper bound on how partial STRIPS model can help on domains that are not modeled completely.

The results support our hypothesis that partial STRIPS model can be useful to enhance the guidance of weak ML-based heuristics and policies. In particular, coverage often increases compared to the variant that uses only the HGN heuristics. There are only two cases, marked in red, where there is a slight decrease in coverage.

Nevertheless, as expected, the quality of the guidance of a partial model heavily depends on the part of the problem that has been modelled. In some cases guidance is as good as with the perfect model. For example, in Grid-keys, by modelling only the position of the keys and ignoring the robots, the search is guided towards relevant sub-goals (picking-up necessary keys and/or dropping them down at their locations). In other cases, however, guidance is more limited. For example, the partial model Logistics-air, ignores completely all locations within a city so it assigns a heuristic value of 0 to all states where the packages are in the right city, regardless of whether the goal has been fulfilled. This causes huge heuristic plateaus and hence the coverage is 0 when used alone. The great news is that, even in such cases, the heuristic can be a good complement for other ML-based heuristics. As shown in the results from Table 1, coverage increases with respect to using only h^{hgn} . A likely reason is that h_A^{FF} helps to guide the search until having all packages in the right city, and then $\pi^{hgn}/\pi^{hgn-one}$ can help to achieve the goal from there.

Figure 1 shows a more detailed view on the comparison on search effort in terms of expansions of only using HGN

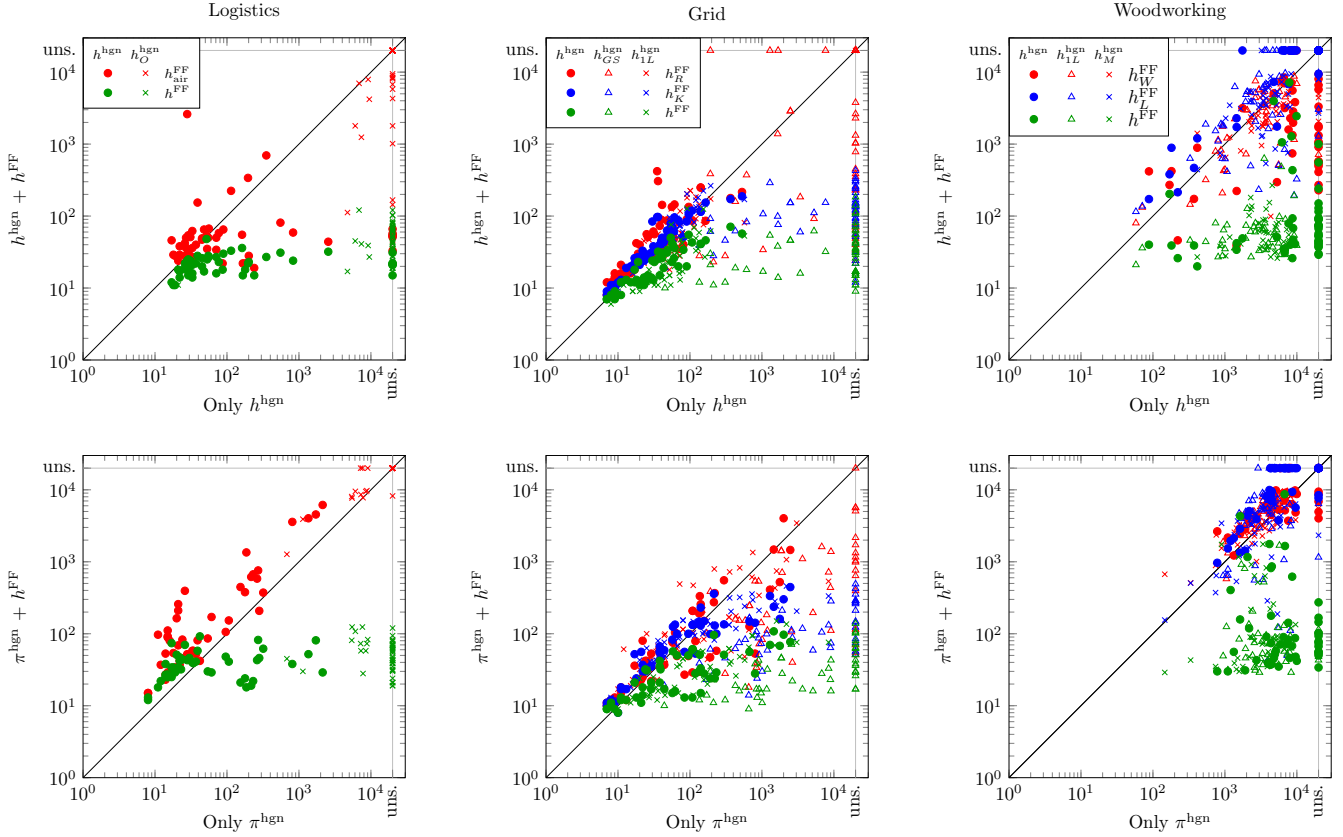


Figure 1: Expansions of only HGN against using both HGN and h^{FF} with different partial models. In the top row, HGN is used as heuristic (e.g. h^{hgn}) and in the bottom row, it is used as a policy (e.g. π^{hgn}). We distinguish different variants of HGN trained with different datasets with shapes, and different partial STRIPS models with colors.

with respect to combining HGN and FF under different partial models. In general, combining HGN and partial STRIPS leads to more points under the diagonal, validating our hypothesis. This is partially due to the way heuristics are combined, by having two separate open lists one for each heuristic, which keeps the overhead small even when the heuristic provided by the partial model is not accurate (e.g. in Grid when using the robot partial model, h_R^{FF} , and the HGN trained on a non-biased dataset). On the other hand, partial models can reduce search effort significantly by several orders of magnitude, whenever they model a part of the domain that was not correctly captured by the HGN heuristic. We remark again that, the computation of the FF heuristic does not have a huge overhead with respect to HGN, and hence the advantages in terms of expanded nodes also carry to search time.

Regarding the cost of the plans found by the different approaches, we observe that there is not a very high variance. As shown in Figure 2, even when using the best HGN variant, plan costs are slightly better when using the partial-model heuristics (finding better plans in 47 tasks compared to 24). When excluding the full model configurations (h^{FF}), this is tied (improving in 19 tasks compared to 18), and if the biased HGN variants are considered the advantage is clearer

(132 tasks compared to 55). Overall, results regarding plan cost seem to correlate with tasks solved, so that more informative heuristics also find (slightly) better plans.

6.2 Other ways of combining heuristics

Table 1 focuses exclusively on the performance using the double-queue algorithm. In general, this was the most suitable way for combining both heuristics. However, it is not always the best. Table 2 shows the cases where the tie-breaking combination mentioned in section 3.3 performed better in the Woodworking-PD domain. We observe that, especially with π^{hgn} , the double-queue algorithm can be outperformed by tie-breaking. Typically, it is better to choose as primary heuristic the one that performs best in isolation, so in this case we used h^{FF} as primary heuristic, breaking ties with h^{hgn} .

Figure 3 shows the cumulative coverage by expanded nodes of the different algorithms. Each plot presents a specific configuration: (a) and (c) shows the performance of h^{hgn} ; (b) and (d) of π^{hgn} . Subfigure 3(a) shows that, combining with the partial model, Woodworking-PD logistics outperforms all other configurations and dramatically increases the number of solved instances. On the other hand, (b) shows that FF tiebreaking with π^{hgn} outperforms a double-queue




Figure 2: Plan cost when using only h^{hgn} against using the double queue method with h^{hgn} and partial-model heuristics, including the h^{FF} .

	none	h_L^{FF}	h_W^{FF}	h^{FF}
none		2	19	50
$h^{\text{hgn-move}}$	14	24+3	24+5	50
h^{hgn}	16	11+4	36+3	50
$h^{\text{hgn-oneloc}}$	26	33+1	40	50
π^{hgn}	10	11+1	14+ 21	49
$\pi^{\text{hgn-oneloc}}$	12	24	10+ 22	50
$\pi^{\text{hgn-move}}$	17	23	18+5	49+1

Table 2: Woodworking-PD. Additional coverage by using an alternative algorithm vs table 1c. See legend in that table. Bold numbers show using using an alternative algorithm lead to a new case where the combination outperforms both HGN and STRIPS models in isolation.

algorithm. In both, using some form of a combination of heuristics is better than using the only one. In the Grid domain, we observe that using h^{hgn} and π^{hgn} , the three configurations that combine the heuristics outperform the use of only one and work similarly.

However, as shown in table 1 and Figure 1, there are situations where the algorithms we considered did not manage to improve over the HGN baseline.

7 Related work

Planning as heuristic search approaches (Bonet and Geffner 2001) use the STRIPS model to relax or transform the state space and guide the search, either by estimating the distance from each state to the goal, or by computing preferences over which actions or state to explore first (Richter and Helmert 2009). These last two ideas are related to the notions used in RL we mention in Section 3.3: value-estimation and action-selection (Silver et al. 2016).

Beyond RL, ML sequence and structured prediction aims to produce sequences with high likelihood. They can be seen as policy estimations as the likelihoods are calculating using an internal state, the past decisions and classifying over

the next possible tokens in the sequence or structure. For instance, (Scholak, Schucher, and Bahdanau 2021) reports using a pre-trained language model to map natural language questions into syntactically correct SQL queries. Typically, beam search is used to find for a sequence with high aggregated likelihood. This case includes the widely used large language models (Vaswani et al. 2017).

There has been work on comparing learned heuristics with classical planning heuristics (Ferber, Helmert, and Hoffmann 2020). It is known that declarative models enable efficient computation of heuristics that are informative and time-efficient (Geffner and Bonet 2013). As such, when a full model is available those heuristics can typically outperform NN heuristics even when training in a fixed task (Ferber, Helmert, and Hoffmann 2020). However, when a full model is not available, it is not clear if and how relaxation-based heuristics can be used.

Previous work has used symbolic information to improve the behaviour of RL algorithms. For instance, Lee et al. (2021) use planning actions to define RL options to be used at training, and Illanes et al. (2020) use symbolic plans for the task to improve the learning stage. Both De Giacomo et al. (2019) and Toro-Icarte et al. (2022) use LTL for specifying RL rewards. On the other hand, Alshiekh et al. (2018) enforce LTL formulas during training and inference so the actions taken are safe. On the planning side, there has been work on obtaining or improving planning heuristics using ML (Yoon, Fern, and Givan 2006; Virseda, Borrajo, and Alcázar 2013; Karia and Srivastava 2021; Ferber et al. 2021; Toyer et al. 2018). Some others have worked on using DRL for generalized planning for solving planning instances in a given domain (Rivlin, Hazan, and Karpas 2020). In contrast with these directions, our work emphasizes modifying the behaviour after the ML model has been trained, amortizing the training cost. Moreover, our contributions suggest new ways for faster adaptation to new specific domains as far as part of the high level structure can be expressed as a (partial) STRIPS model.

Modelling planning problems is studied in knowledge engineering (Vaquero et al. 2013). Recent efforts have looked at obtaining planning models from source code using annotations (Katz, Moshkovich, and Karpas 2018).

Model-lite planning approaches (Kambhampati 2007; Weber and Bryce 2011; Zhuo and Kambhampati 2017) have considered incomplete planning models in the past, e.g., for the generation of robust plans. Instead, we explore how to use them to enhance existing ML-based systems.

8 Discussion

Pure RL methods can generalize well when the training data covers well-enough the distribution of trajectories (Sutton and Barto 2018). On the other hand, it is common to find that RL models show weak generalization when they rely on passively collected data. In this case, the most common next steps are tuning the ML algorithm or collecting more data based on insights of error analysis. In this work we explored a complementary direction: improve the scalability by turning such insights into a symbolic model, exploiting classical planners which naturally deal well with gen-

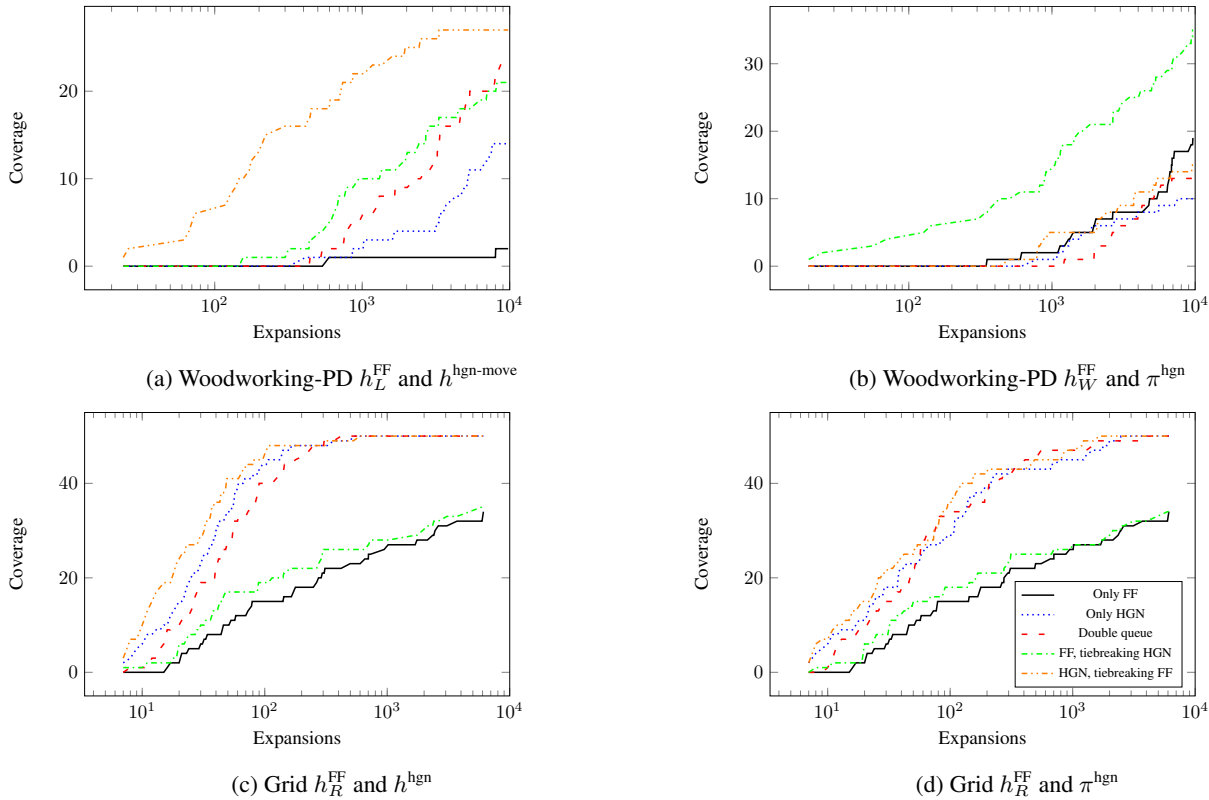


Figure 3: Cumulative coverage by the number of expanded nodes under different variants of h^{FF} , h^{hgn} and π^{hgn} .

eralization. We showed that, even when a full model is not available, classical planning heuristics can complement well ML-based heuristics. Typically, more accurate models lead to better search guidance. But even models that disregard large parts of the problem can be very beneficial.

In general, the actions returned by a classical planner can be executed safely and reach the goal as far as the model correctly represents the real state space. However, classical planners are also used in cases where this not the case. For instance, simple classes of full-observable non-deterministic planning can be tackled by assuming a deterministic problem and replanning when the execution fails (Yoon, Fern, and Givan 2007). While this work is about black-box planning, it can also be seen as using planners while associating planning states with external states. Complementing ML-based methods with model-based heuristics could improve the impact of planning research, and lead to a deeper understanding of the relationship between learning and reasoning. For instance, one could start with a trained policy where states are partially symbolic, and associate that part of the state to an STRIPS model. In turn, this combination could offer an stronger baseline for further research on DRL.

As future work, there are many directions worth pursuing. We would like to test our methods using other ML models and in other domains. For instance, our approach would be directly applicable on domains where the dynamics of the black-box simulator cannot be expressed in STRIPS, e.g., with non-symbolic states or continuous variables.

Another question is what are good partial STRIPS models. In some cases, the ML models and the partial STRIPS model might be focusing on different aspects of the problem, which could be in contradiction. For instance, as training ML aims to minimize the expected cost, we might want to use a STRIPS model to avoid taking actions leading to undesirable states, leading the search towards more secure zones. In principle, this would make the plans more reliable but we have not explored the overall behaviour of our algorithms in such cases. One way to study this issue is to consider ML models trained on data about less risky scenarios, while the partial STRIPS model focus on safety issues.

Finally, our tie-breaking configurations commit exclusively into the preferences of the first heuristic. Hence, it would interesting to explore other ways of balancing both heuristics, such as focal-search based algorithms (Greco and Baier 2021; Greco, Araneda, and Baier 2022; Greco et al. 2022a).

Acknowledgements

We would like to thank William Shen and Florian Geißer for sharing the code of their implementation of the hypergraph networks (HGN) in Fast Downward. Matias Greco was supported by the National Agency for Research and Development (ANID) / Doctorado Nacional / 2019 - 21192036. Matias Greco and Jorge Baier thank to Centro Nacional de Inteligencia Artificial CENIA, FB210017, BASAL, ANID.

References

- Alshiekh, M.; Bloem, R.; Ehlers, R.; Könighofer, B.; Niekum, S.; and Topcu, U. 2018. Safe reinforcement learning via shielding. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*, volume 32.
- Araneda, P.; Greco, M.; and Baier, J. A. 2021. Exploiting Learned Policies in Focal Search. In Ma, H.; and Serina, I., eds., *Proceedings of the 14th Annual Symposium on Combinatorial Search (SoCS 2021)*, volume 12, 2–10. AAAI Press.
- Arulkumaran, K.; Deisenroth, M. P.; Brundage, M.; and Bharath, A. A. 2017. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6): 26–38.
- Bäckström, C.; and Nebel, B. 1995. Complexity Results for SAS⁺ Planning. *Computational Intelligence*, 11(4): 625–655.
- Bonet, B.; and Geffner, H. 2001. Planning as Heuristic Search. *Artificial Intelligence*, 129(1): 5–33.
- Culberson, J. C.; and Schaeffer, J. 1998. Pattern Databases. *Computational Intelligence*, 14(3): 318–334.
- De Giacomo, G.; Iocchi, L.; Favorito, M.; and Patrizi, F. 2019. Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, 128–136.
- 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.
- Ferber, P.; Geißer, F.; Trevizan, F.; Helmert, M.; and Hoffmann, J. 2021. Neural Network Heuristic Functions for Classical Planning: Reinforcement Learning and Comparison to Other Methods. In Palacios, H.; Gómez, V.; Sanner, S.; Jonsson, A.; Kolobov, A.; and Fern, A., eds., *ICAPS 2021 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL)*.
- Ferber, P.; Helmert, M.; and Hoffmann, J. 2020. Neural Network Heuristics for Classical Planning: A Study of Hyperparameter Space. In De Giacomo, G., ed., *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, 2346–2353. IOS 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.
- Geffner, H.; and Bonet, B. 2013. *A Concise Introduction to Models and Methods for Automated Planning*, volume 7 of *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool.
- Greco, M.; Araneda, P.; and Baier, J. A. 2022. Focal Discrepancy Search for Learned Heuristics. In *Proceedings of the 15th Annual Symposium on Combinatorial Search (SoCS 2022)*, volume 13.
- Greco, M.; and Baier, J. A. 2021. Bounded-Suboptimal Search with Learned Heuristics. *ICAPS Workshop PRL*.
- Greco, M.; Toro, J.; Hernandez, C.; and Baier, J. A. 2022a. K-Focal Search for Slow Learned Heuristics. In *Proceedings of the 15th Annual Symposium on Combinatorial Search (SoCS 2022)*, volume 13.
- Greco, M.; Álvaro Torralba; Baier, J. A.; and Palacios, H. 2022b. Code and data for the RDDPS 2022 paper “Scaling up ML-based Black-box Planning with Partial STRIPS Models”. <https://doi.org/10.5281/zenodo.6544069>.
- Harvey, W. D.; and Ginsberg, M. L. 1995. Limited Discrepancy Search. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI 1995)*, 607–615. Morgan Kaufmann.
- Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research*, 26: 191–246.
- Hoffmann, J. 2005. Where ‘Ignoring Delete Lists’ Works: Local Search Topology in Planning Benchmarks. *Journal of Artificial Intelligence Research*, 24: 685–758.
- Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14: 253–302.
- Illanes, L.; Yan, X.; Icarte, R. T.; and McIlraith, S. A. 2020. Symbolic plans as high-level instructions for reinforcement learning. In Beck, J. C.; Karpas, E.; and Sohrabi, S., eds., *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)*, 540–550. AAAI Press.
- Kambhampati, S. 2007. Model-lite Planning for the Web Age Masses: The Challenges of Planning with Incomplete and Evolving Domain Models. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI 2007)*, 1601–1605. AAAI Press.
- Karia, R.; and Srivastava, S. 2021. Learning Generalized Relational Heuristic Networks for Model-Agnostic Planning. In Leyton-Brown, K.; and Mausam, eds., *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*, 8064–8073. AAAI Press.
- Karoui, W.; Huguot, M.-J.; Lopez, P.; and Naanaa, W. 2007. YIELDS: A yet improved limited discrepancy search for CSPs. In *CPAIOR*, 99–111. Springer.
- Katz, M.; Moshkovich, D.; and Karpas, E. 2018. Semi-Black Box: Rapid Development of Planning Based Solutions. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*, volume 32, 6211–6218.
- Lee, J.; Katz, M.; Agravante, D. J.; Liu, M.; Klinger, T.; Campbell, M.; Sohrabi, S.; and Tesauro, G. 2021. AI Planning Annotation in Reinforcement Learning: Options and Beyond. In Palacios, H.; Gómez, V.; Sanner, S.; Jonsson, A.; Kolobov, A.; and Fern, A., eds., *ICAPS 2021 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL)*.
- McDermott, D. 2000. The 1998 AI Planning Systems Competition. *AI Magazine*, 21(2): 35–55.
- McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. 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.
- Richter, S.; and Helmert, M. 2009. Preferred Operators and Deferred Evaluation in Satisficing Planning. In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS 2009)*, 273–280. AAAI Press.
- Rivlin, O.; Hazan, T.; and Karpas, E. 2020. Generalized Planning With Deep Reinforcement Learning. In Katz, M.; Palacios, H.; Fern, A.; Gómez, V.; Jonsson, A.; and Sanner, S., eds., *ICAPS 2020 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL)*, 16–24.
- Scholak, T.; Schucher, N.; and Bahdanau, D. 2021. PICARD: Parsing Incrementally for Constrained Auto-Regressive Decoding from Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 9895–9901.
- Seipp, J.; Pommerening, F.; Sievers, S.; and Helmert, M. 2017. Downward Lab. <https://doi.org/10.5281/zenodo.790461>.
- Shen, W.; Trevizan, F.; and Thiébaux, S. 2020. Learning Domain-Independent Planning Heuristics with Hypergraph Networks. In Beck, J. C.; Karpas, E.; and Sohrabi, S., eds., *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)*, 574–584. AAAI Press.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature*, 529(7587): 484–489.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press.
- Toro-Icarte, R.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2022. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research*, 73: 173–208.
- Toyer, S.; Trevizan, F.; Thiébaux, S.; and Xie, L. 2018. Action Schema Networks: Generalised Policies with Deep Learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*, 6294–6301.
- Vaquero, T. S.; Silva, J. R.; Tonidandel, F.; and Beck, J. C. 2013. itSIMPLE: towards an integrated design system for real planning applications. *The Knowledge Engineering Review*, 28(2): 215–230.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need.
- Virseda, J.; Borrajo, D.; and Alcázar, V. 2013. Learning heuristic functions for cost-based planning. In *Proceedings of ICAPS workshop on Planning and Learning*.
- Weber, C.; and Bryce, D. 2011. Planning and acting in incomplete domains. 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)*, 274–281. AAAI Press.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3): 229–256.
- Yoon, S.; Fern, A.; and Givan, R. 2006. Learning Heuristic Functions from Relaxed Plans. In Long, D.; Smith, S. F.; Borrajo, D.; and McCluskey, L., eds., *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling (ICAPS 2006)*, 162–171. AAAI Press.
- Yoon, S.; Fern, A.; and Givan, R. 2007. FF-Replan: A Baseline for Probabilistic Planning. In Boddy, M.; Fox, M.; and Thiébaux, S., eds., *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling (ICAPS 2007)*, 352–360. AAAI Press.
- Zhuo, H. H.; and Kambhampati, S. 2017. Model-lite planning: Case-based vs. model-based approaches. *Artificial Intelligence*, 246: 1–21.