
Abstract Interpretation for Generalized Heuristic Search in Model-Based Planning

Tan Zhi-Xuan¹ Joshua B. Tenenbaum¹ Vikash K. Mansinghka¹

Abstract

Domain-general model-based planners often derive their generality by constructing search heuristics through the relaxation or abstraction of symbolic world models. We illustrate how abstract interpretation can serve as a unifying framework for these abstraction-based heuristics, extending the reach of heuristic search to richer world models that make use of more complex datatypes and functions (e.g. sets, geometry), and even models with uncertainty and probabilistic effects.

1. Introduction

Since the advent of the Stanford Research Institute Problem Solver (STRIPS) (Fikes & Nilsson, 1971), it has been understood that planning and sequential decision making can be viewed as a form of theorem proving, where sequences of actions are derived via heuristic search given a symbolic world model and goal. While similar formal approaches have subsequently been highly successful in model checking (Clarke, 1997), program analysis (Nielson et al., 2004), and constraint satisfaction (Barrett & Tinelli, 2018), they have come to play less of a role in automated planning and decision making, in light of successful learning-based approaches. (Mnih et al., 2015; Silver et al., 2017). This is in part due to the difficulty of specifying accurate symbolic models of the world, given the presence of uncertainty and limited expressiveness, e.g., to models with only propositional variables (Jiménez et al., 2012). Nonetheless, symbolic methods remain widely used in planning and robotics, often in conjunction with machine learning or motion planning (Leonetti et al., 2016; Garrett et al., 2021).

Here, we show one powerful way of making these symbolic planning methods more general: Using abstract interpretation (Cousot & Cousot, 1977) to construct generalized heuristics for search. Domain-general planners often construct search heuristics by planning in a *relaxed* or *abstracted* model: the cost of a solution in a relaxed model

can be used as an (optimistic) estimate of the true cost, providing heuristic guidance in search algorithms. Some of the abstractions used by these heuristics are also used in model checking (Seipp & Helmert, 2018) or numeric invariant analysis (Scala et al., 2016). However, they have typically been limited to propositional variables, with a few numeric extensions. Inspired by work on semantic modularity for symbolic planners (Gregory et al., 2012), we illustrate how abstract interpretation can serve as a unifying framework for these abstraction-based heuristics. This leads to natural extensions of heuristic search to richer world models defined using more complex datatypes (e.g. sets) and functions (e.g. geometric operations), or even models with uncertainty and probabilistic effects. These heuristics can also be integrated with learning, allowing agents to jumpstart planning in novel world models via abstraction-derived information that is later refined by experience. This suggests that abstract interpretation can play a key role in building universal reasoning systems.

2. Planning in Symbolic World Models

Symbolic world models, also called *planning domains*, describe the transition dynamics of an environment in symbolic terms. A planning domain constitutes part of a *planning problem*, which also specifies the initial state and goal. While domains and problems are usually specified in a first-order language such as PDDL (McDermott et al., 1998), we adopt a simplified formalism for ease of exposition. Here, a planning problem is a tuple (F, V, A, I, G) , where:

- F is a set of functions;
- V is a set of (typed) variables which may have Boolean, numeric, or other datatypes;
- A is a set of actions $a \in A$, where:
 - $\text{pre}(a)$ is a precondition for a , comprising Boolean functions from F over variables in V ;
 - $\text{eff}(a)$ is the effect of a , a set of non-conflicting assignments to variables in $v \in V$
- I is the initial state, a set of assignments to each $v \in V$;
- G is the goal, a logical formula of Boolean functions from F over variables in V

¹Massachusetts Institute of Technology. Correspondence to: Tan Zhi-Xuan <xuan@mit.edu>.

```

board.person1.plane1.city1
pre: person1.loc = city1.loc ∧ plane1.loc = city1.loc
eff: person1.loc := plane1.loc;
    onboard1 := add.elem(onboard1, person1.id)
fly.city1.city2
pre: plane1.loc = city1.loc ∧
    fuel1 > burn.rate * cardinality(onboard1)
eff: plane1.loc := city2.loc;
    fuel1 := fuel1 - burn.rate * cardinality(onboard1);
    
```

Figure 1. Example actions in an air-travel problem, with finite-domain, numeric, and set-valued variables.

2.1. Action Semantics

Assignment can be assumed to follow the standard semantics for imperative languages. Predicates and functions used in planning domains typically include testing if a Boolean variable is false or true, along with arithmetic operations (Fox & Long, 2003), but can be defined for new datatypes as necessary (Gregory et al., 2012). Since assignments in the effect $\text{eff}(a)$ of an action a must be non-conflicting, $\text{eff}(a)$ corresponds to parallel execution of each assignment. We can hence specify a denotational semantics for a by identifying it with a relation over states $(S, T) \in a$, where $\text{pre}(a) \models S$, and T is the result of applying $\text{eff}(a)$ to S . This is equivalent to the semantics for atomic actions in concurrent programs by Lamport (1990). Alternatively, actions can be seen as (guarded) primitive commands in a domain specific imperative language, with semantics defined by condition-checking and assignment as sub-primitives. Two example actions are shown in Figure 1.

2.2. Planning Algorithms

Common approaches to solving symbolic planning problems include backward search from the goal G , or forward best-first search from the initial state I (Bonet & Geffner, 2001), with the fastest planners typically relying on variants of forward A* search guided by highly informative heuristics (Helmert, 2006), along with compilation techniques for speed and memory efficiency (Helmert, 2009; Zhi-Xuan, 2022). Heuristic-guided search can also be extended to probabilistic and partially observable domains, using algorithms such as Real Time Dynamic Programming (RTDP) (Bonet & Geffner, 2003) and Trial-based Heuristic Tree Search (THTS) (Keller & Helmert, 2013).

3. Abstract Semantics for Symbolic Planning

In order to perform abstract interpretation over symbolic world models, we need an abstraction function $\alpha(\cdot)$ that maps a concrete set of states \mathcal{S} to an abstract state \mathcal{S}^\sharp , along with a concretization function $\gamma(\cdot)$ s.t. $\mathcal{S} \subseteq \gamma(\alpha(\mathcal{S}))$, where each concrete state $S \in \mathcal{S}$ is a complete assignment of values to the variables V of the planning problem. In addition,

we have to provide abstract semantics for each action, effectively abstract actions a^\sharp that define a relation over abstract states $(\mathcal{S}^\sharp, \mathcal{T}^\sharp) \in a^\sharp$. We now describe several abstractions that are useful for deriving heuristics.

3.1. State Abstractions

Cartesian Abstraction. One approach to abstracting states is to assign an abstract value to each variable $v \in V$, representing a set of concrete values v could take. For example, the numeric variable `fuel1` in Figure 1 could be assigned the interval value $[0, 1000)$. Doing this for each variable produces a *Cartesian abstraction*: Given the values in an abstract state \mathcal{S}^\sharp , the corresponding set of concrete states $\gamma(\mathcal{S}^\sharp)$ is just the Cartesian product of sets associated with each abstract value (Ball et al., 2001).

Predicate Abstraction. Another approach is to consider the set of all predicates \mathcal{P} that feature in action preconditions $\text{pre}(a)$ for all $a \in A$, or the goal specification G . We define abstract states to be sets of these predicates $\mathcal{S}^\sharp \subseteq \mathcal{P}$, corresponding to the set of concrete states where every predicate in the abstract state holds true. The advantage of this abstraction is that it can be easily used to determine if an action is executable, or if the goal is achieved.

3.2. Action Abstractions

State-Induced. Given abstraction and concretization functions α and γ , this induces a corresponding abstract action a^\sharp for each action a : Given an abstract state \mathcal{S}^\sharp , we have $(\mathcal{S}^\sharp, \mathcal{T}^\sharp) \in a^\sharp$ where $\mathcal{T}^\sharp := \alpha(\{T \mid S \in \gamma(\mathcal{S}^\sharp), (S, T) \in a\})$. This can be efficiently implemented for Cartesian abstractions when the right-hand-side of all assignments in an action is constant, since the result of applying such an action is a specific concrete state which can then be abstracted.

Widening-Based. Given a Cartesian state abstraction, a more relaxed abstraction is to replace each assignment $v := E$ in a concrete action a with $v := v \nabla E$, where E is an expression and ∇ is a choice of widening operator associated with the value abstraction for variable v (Gregory et al., 2012). In the case where v is Boolean or has a finite domain, a natural abstraction for v is to replace the domain of v with its powerset (e.g., allowing a Boolean v to be true and false at the same time), using the join operator \sqcup as a widening. Other choices are necessary when v is numeric and unbounded, e.g. delayed widening (Miné, 2017).

Widening-based abstract actions can be directly applied to Cartesian state abstractions, since they modify each state variable independently. They can also be extended to combined Cartesian and predicate abstractions, where each abstract state $\mathcal{S}^\sharp = (X^\sharp, P)$ comprises both abstract assignments X^\sharp and predicates P . The successor of applying some action a^\sharp is then $(\nabla_a X^\sharp, Q)$, where $\nabla_a X^\sharp$ is the widened




Figure 2. A graphical analogy between (a) abstract interpretation for error detection and (b) abstract planning for goal reachability. In abstract interpretation, over-approximating the set of reachable states via abstract steps allows us to rule out impossible errors. In abstract planning, we can similarly rule out some goals as unreachable, or otherwise estimate how many steps it takes to reach them.

set of assignments, and Q is the largest set of predicates in \mathcal{P} which hold true for at least some concrete $s \in \gamma(\nabla_a X^\sharp)$. Note that $P \subseteq Q$ by construction (widening can only cause more predicates to be true), and Q can be computed efficiently from P and $\text{eff}(a^\sharp)$ by checking and adding only the predicates in \mathcal{P} that are affected by $\text{eff}(a^\sharp)$.

4. Deriving Heuristics via Abstraction

Having described several abstractions for symbolic world models, we now illustrate how they can be used to derive multiple families of planning heuristics developed in the literature. We discuss novel extensions along the way.

4.1. Relaxed Reachability via Widening

As shown in Figure 2, determining the reachability of a goal condition in symbolic planning can be viewed as analogous to determining the reachability of an error in program verification. This analysis can be performed by imagining a non-deterministic search procedure that repeatedly applies all executable actions at once, then over-approximating that search process by merging the results of all actions. As noted by Gregory et al. (2012), this is equivalent to repeatedly applying widened versions of the actions until a goal condition or fixpoint is reached (though they do not describe it in terms of abstract interpretation). The number of iterations required serves as a lower bound on the true number of steps to the goal, and hence an optimistic heuristic estimate that can guide search. Indeed, when all variables are propositional, this is equivalent to the h_{\max} delete relaxation heuristic introduced by Bonet & Geffner (2001).

One issue with this analysis is that it may not terminate if variables have infinite domains and the widening operator ∇ is not appropriately chosen. To address this, Gregory et al. (2012) employ an equivalent of delayed widening. However, many other widening strategies have been used in abstract interpretation, such as widening with thresholds (Miné, 2017), that could lead to improved trade-offs between informativeness and convergence time when computing the heuristic.

4.2. Relaxed Reachability over Predicate Sets

An alternative formulation of the widening-based reachability analysis is to consider planning over *predicate sets* using the predicate abstraction described earlier. More specifically, consider the (abstract) state space graph induced by the combined Cartesian and predicate abstraction, with states $S^\sharp = (X^\sharp, P)$. Following Haslum & Geffner (2000), finding the shortest path to a goal-satisfying state in the original problem is equivalent to finding the shortest path to an abstract state (X^\sharp, P) where $P = G$ (we assume w.l.o.g. that the goal G is a conjunction of predicates).

Now we perform several relaxations on this state space graph. First, we add edges corresponding to the widened actions described earlier, such that (X^\sharp, P) connects to $(\nabla_a X^\sharp, Q)$ for a widened action a^\sharp . Next, for every node (X^\sharp, P) that has $(\nabla_a X^\sharp, Q)$ as a successor, we add an edge from (X^\sharp, P) to each node (\cdot, R) where $Q \rightarrow R$. Finally, for each node (X^\sharp, Q) , we add a zero-cost edge *from* every single-predicate node (\cdot, R) s.t. $R \subseteq Q$ and $|R| = 1$. These relaxations achieve the following facts:

- The cost of the shortest path in the relaxed graph C is now less than the cost of the shortest path in the original problem C^* .
- By the second relaxation, we can reach every single-predicate node (\cdot, R) , $|R| = 1$ with less than or equal cost to reaching a multi-predicate node.
- By the third relaxation, we can compute the cost of reaching any multi-predicate node $C(\cdot, Q)$ as the max over costs of achieving each predicate in that node:

$$C(\cdot, Q) = \max_{R \subseteq Q, |R|=1} C(\cdot, R)$$

When all variables V are propositional, and the right hand sides of all assignments are constant, this again recovers the h_{\max} heuristic, reformulated as the h^1 heuristic in Haslum & Geffner (2000). However, by incorporating the widening-based abstraction from earlier, this generalizes the heuristic to numeric variables, or any other datatype with an effective abstraction (e.g. approximating geometric sets with zonotopes (Bogomolov et al., 2019)). Furthermore, this derivation allows for a number of simple modifications, such as computing the cost of achieving a multi-predicate node as the *sum* instead of the maximum over the costs of each individual predicate. This gives a generalized version of the widely used additive heuristic h_{add} , subsuming existing work on numeric subgoaling heuristics (Scala et al., 2020). While the result heuristic is non-admissible, it is generally much more informative and useful for search.

4.3. Counterexample Guided Abstraction Refinement

A final class of abstraction heuristics enabled by abstract interpretation is counter-example guided abstraction refinement (CEGAR) (Seipp & Helmert, 2018). Starting with a coarse-grained Cartesian abstraction, one can solve for abstract plans to estimate the distance to the goal, and then iteratively refine the abstraction when counterexamples to the correctness of the abstraction are found. Each time after refinement, heuristic values obtained from the previous iteration can be used to guide search in the current iteration, leading to a rapid iterative approach that automatically constructs useful abstractions for planning.

However, CEGAR heuristics have so far been limited to problems with propositional variables, and CEGAR-like algorithms have only recently been explored in symbolic-geometric contexts for task and motion planning (Thomason & Kress-Gazit, 2021). By leveraging abstract semantics introduced for other datatypes in program analysis, it may be possible to extend the reach of such heuristics to a broader array of problems and domains.

5. Other Extensions and Applications

Thus far, we have focused on how abstract interpretation can be used to derive more general heuristics for forward search. However, there are many other possible applications of abstract interpretation, which we briefly reflect upon here.

5.1. Uncertainty and Learning

As noted earlier, heuristics can be used within algorithms such as Real Time Dynamic Programming (RTDP) (Bonet & Geffner, 2003) and Trial-based Heuristic Tree Search (THTS) (Keller & Helmert, 2013) that operate over stochastic domains, using them to initialize an estimate of the value function for the underlying Markov Decision Process. To derive these heuristics from formal analysis of domains with probabilistic effects, the simplest abstraction would be to assume that all branches of a (discrete) probabilistic choice occur at once. This would effectively relax the problem, enabling the computation of admissible heuristics that guarantee eventual convergence to the true value function (Barto et al., 1995). For continuous probability distributions with bounded support, it would also be possible to abstract distributions with their upper and lower bounds.

Using abstraction heuristics to initialize value functions is closely related to the idea of using heuristics as a learning target for neural network estimators of expected value (Shen et al., 2020; Gehring et al., 2021). This would allow agents to jumpstart planning in novel world models via abstraction-derived information that is later refined by experience.

5.2. Reverse Interpretation for Bidirectional Search

Abstract interpretation can be executed in reverse (Hughes, 1987; Monniaux, 2001), allowing for generalizations of backward search and heuristic construction to domains with non-propositional variables. This might allow for novel combinations of (abstract) backwards search with (concrete) forward search, as has been explored in robotics (Garrett et al., 2018) and program analysis (Dinges & Agha, 2014).

5.3. Abstract Interpretation for Generalized Planning

Generalized plans, or policies, often take the form of imperative programs with control flow (Andre & Russell, 2002; Jiménez et al., 2019; Segovia-Aguas et al., 2021). This suggests that abstraction-guided program synthesis methods (Solar-Lezama, 2008; Srivastava et al., 2010; Wang et al., 2017) can also be used for generalized planning.

This extended abstract presents only an initial foray into the manifold connections between abstract interpretation (AI) and artificially intelligent (AI) planning. We hope that illustrating some of these connections lays the ground for future interdisciplinary work.

References

- Andre, D. and Russell, S. J. State abstraction for programmable reinforcement learning agents. In *Aaai/iaai*, pp. 119–125, 2002.
- Ball, T., Podelski, A., and Rajamani, S. K. Boolean and cartesian abstraction for model checking c programs. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pp. 268–283. Springer, 2001.
- Barrett, C. and Tinelli, C. Satisfiability modulo theories. In *Handbook of model checking*, pp. 305–343. Springer, 2018.
- Barto, A. G., Bradtke, S. J., and Singh, S. P. Learning to act using real-time dynamic programming. *Artificial intelligence*, 72(1-2):81–138, 1995.
- Bogomolov, S., Forets, M., Frehse, G., Potomkin, K., and Schilling, C. Juliareach: a toolbox for set-based reachability. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, pp. 39–44, 2019.
- Bonet, B. and Geffner, H. Planning as heuristic search. *Artificial Intelligence*, 129(1-2):5–33, 2001.
- Bonet, B. and Geffner, H. Labeled RTDP: Improving the convergence of real-time dynamic programming. In *ICAPS*, volume 3, pp. 12–21, 2003.
- Clarke, E. M. Model checking. In *International Conference on Foundations of Software Technology and Theoretical Computer Science*, pp. 54–56. Springer, 1997.
- Cousot, P. and Cousot, R. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pp. 238–252, 1977.
- Dinges, P. and Agha, G. Targeted test input generation using symbolic-concrete backward execution. In *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering*, pp. 31–36, 2014.
- Fikes, R. E. and Nilsson, N. J. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4):189–208, 1971.
- Fox, M. and Long, D. PDDL2. 1: An extension to PDDL for expressing temporal planning domains. *Journal of artificial intelligence research*, 20:61–124, 2003.
- Garrett, C. R., Lozano-Perez, T., and Kaelbling, L. P. Ffrob: Leveraging symbolic planning for efficient task and motion planning. *The International Journal of Robotics Research*, 37(1):104–136, 2018.
- Garrett, C. R., Chitnis, R., Holladay, R., Kim, B., Silver, T., Kaelbling, L. P., and Lozano-Pérez, T. Integrated task and motion planning. *Annual review of control, robotics, and autonomous systems*, 4:265–293, 2021.
- Gehring, C., Asai, M., Chitnis, R., Silver, T., Kaelbling, L. P., Sohrabi, S., and Katz, M. Reinforcement learning for classical planning: Viewing heuristics as dense reward generators. *arXiv preprint arXiv:2109.14830*, 2021.
- Gregory, P., Long, D., Fox, M., and Beck, J. C. Planning modulo theories: Extending the planning paradigm. In *Twenty-Second International Conference on Automated Planning and Scheduling*, 2012.
- Haslum, P. and Geffner, H. Admissible heuristics for optimal planning. In *The Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS), Breckenridge, Colorado, 14-17 April, 2000.*, pp. 140–149. AAAI Press, 2000.
- Helmert, M. The Fast Downward planning system. *Journal of Artificial Intelligence Research*, 26:191–246, 2006.
- Helmert, M. Concise finite-domain representations for PDDL planning tasks. *Artificial Intelligence*, 173(5-6): 503–535, 2009.
- Hughes, J. *Backwards analysis of functional programs*. University of Glasgow. Department of Computing Science, 1987.
- Jiménez, S., De La Rosa, T., Fernández, S., Fernández, F., and Borrajo, D. A review of machine learning for automated planning. *The Knowledge Engineering Review*, 27(4):433–467, 2012.
- Jiménez, S., Segovia-Aguas, J., and Jonsson, A. A review of generalized planning. *The Knowledge Engineering Review*, 34, 2019.
- Keller, T. and Helmert, M. Trial-based heuristic tree search for finite horizon MDPs. In *Twenty-Third International Conference on Automated Planning and Scheduling*, 2013.
- Lamport, L. win and sin: Predicate transformers for concurrency. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 12(3):396–428, 1990.
- Leonetti, M., Iocchi, L., and Stone, P. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. *Artificial Intelligence*, 241:103–130, 2016.
- McDermott, D., Ghallab, M., Howe, A., Knoblock, C., Ram, A., Veloso, M., Weld, D., and Wilkins, D. PDDL - the Planning Domain Definition Language, 1998.

- Miné, A. Tutorial on static inference of numeric invariants by abstract interpretation. *Foundations and Trends in Programming Languages*, 4(3-4):120–372, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- Monniaux, D. Backwards abstract interpretation of probabilistic programs. In *European Symposium on Programming*, pp. 367–382. Springer, 2001.
- Nielson, F., Nielson, H. R., and Hankin, C. *Principles of program analysis*. Springer Science & Business Media, 2004.
- Scala, E., Haslum, P., Thiébaux, S., and Ramirez, M. Interval-based relaxation for general numeric planning. In *ECAI 2016*, pp. 655–663. IOS Press, 2016.
- Scala, E., Haslum, P., Thiébaux, S., and Ramirez, M. Subgoalting techniques for satisficing and optimal numeric planning. *Journal of Artificial Intelligence Research*, 68: 691–752, 2020.
- Segovia-Aguas, J., Jiménez, S., and Jonsson, A. Generalized planning as heuristic search. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 31, pp. 569–577, 2021.
- Seipp, J. and Helmert, M. Counterexample-guided cartesian abstraction refinement for classical planning. *Journal of Artificial Intelligence Research*, 62:535–577, 2018.
- Shen, W., Trevizan, F., and Thiébaux, S. Learning domain-independent planning heuristics with hypergraph networks. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 30, pp. 574–584, 2020.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Solar-Lezama, A. *Program synthesis by sketching*. University of California, Berkeley, 2008.
- Srivastava, S., Gulwani, S., and Foster, J. S. From program verification to program synthesis. In *Proceedings of the 37th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pp. 313–326, 2010.
- Thomason, W. and Kress-Gazit, H. Counterexample-guided repair for symbolic-geometric action abstractions. *arXiv preprint arXiv:2105.06537*, 2021.
- Wang, X., Dillig, I., and Singh, R. Program synthesis using abstraction refinement. *Proceedings of the ACM on Programming Languages*, 2(POPL):1–30, 2017.
- Zhi-Xuan, T. PDDL.jl: An extensible interpreter and compiler interface for fast and flexible AI planning. Master’s thesis, MIT, 2022.