

Dissertation zur Erlangung des Doktorgrades der Technischen Fakultät der
Albert-Ludwigs-Universität Freiburg im Breisgau

Symbolic Search for Optimal Planning with Expressive Extensions

David Speck



2022

Dean:

Prof. Dr. Roland Zengerle, *University of Freiburg, Germany*

PhD advisor and first reviewer:

Prof. Dr. Bernhard Nebel, *University of Freiburg, Germany*

Second reviewer:

Prof. Dr. Álvaro Torralba, *Aalborg University, Denmark*

Date of defense:

February 23, 2022

Abstract

In classical planning, the goal is to derive a course of actions that allows an intelligent agent to move from any situation it finds itself in to one that satisfies its goals. Classical planning is considered domain-independent, i.e., it is not limited to a particular application and can be used to solve different types of reasoning problems. In practice, however, some properties of a planning problem at hand require an expressive extension of the standard classical planning formalism to capture and model them. Although the importance of many of these extensions is well known, most planners, especially optimal planners, do not support these extended planning formalisms. The lack of support not only limits the use of these planners for certain problems, but even if it is possible to model the problems without these extensions, it often leads to increased effort in modeling or makes modeling practically impossible as the required problem encoding size increases exponentially.

In this thesis, we propose to use symbolic search for cost-optimal planning for different expressive extensions of classical planning, all capturing different aspects of the problem. In particular, we study planning with axioms, planning with state-dependent action costs, oversubscription planning, and top- k planning. For all formalisms, we present complexity and compilability results, highlighting that it is desirable and even necessary to natively support the corresponding features. We analyze symbolic heuristic search and show that the search performance does not always benefit from the use of a heuristic and that the search performance can exponentially deteriorate even under the best possible circumstances, namely the perfect heuristic. This reinforces that symbolic blind search is the dominant symbolic search strategy nowadays, on par with other state-of-the-art cost-optimal planning strategies. Based on this observation and the lack of good heuristics for planning formalisms with expressive extensions, symbolic search turns out to be a strong approach. We introduce symbolic search to support each of the formalisms individually and in combination, resulting in optimal, sound, and complete planning algorithms that empirically compare favorably with other approaches.

Zusammenfassung

Bei der klassischen Planung besteht das Ziel darin, einen Handlungsablauf zu finden, der es einem intelligenten Agenten ermöglicht, aus jeder Situation, in der er sich befindet, in eine Situation zu gelangen, die seine Ziele erfüllt. Die klassische Planung gilt als domänenunabhängig, d.h. sie ist nicht auf eine bestimmte Anwendung beschränkt und kann zur Lösung verschiedener Arten von Logikproblemen verwendet werden. In der Praxis erfordern jedoch einige Eigenschaften eines vorliegenden Planungsproblems eine ausdrucksstarke Erweiterung des klassischen Standardplanungsformalismus, um sie zu erfassen und zu modellieren. Obwohl die Bedeutung vieler dieser Erweiterungen bekannt ist, unterstützen die meisten Planer, insbesondere optimale Planer, diese erweiterten Planungsformalismen nicht. Die fehlende Unterstützung schränkt nicht nur die Verwendung dieser Planer für bestimmte Probleme ein, sondern selbst wenn es möglich ist, die Probleme ohne diese Erweiterungen zu modellieren, führt dies oft zu einem erhöhten Aufwand bei der Modellierung oder macht die Modellierung praktisch unmöglich, da die erforderliche Problemkodierungsgröße exponentiell ansteigt.

In dieser Arbeit schlagen wir vor, die symbolische Suche für kostenoptimale Planung für verschiedene ausdrucksstarke Erweiterungen der klassischen Planung zu verwenden, die alle unterschiedliche Aspekte des Problems erfassen. Insbesondere untersuchen wir die Planung mit Axiomen, die Planung mit zustandsabhängigen Aktionskosten, die "Oversubscription Planung" und die "Top- k Planung". Für alle Formalismen präsentieren wir Ergebnisse zur Komplexität und Kompilierbarkeit, wobei wir hervorheben, dass es wünschenswert und sogar notwendig ist, die entsprechenden Aspekte nativ zu unterstützen. Wir analysieren die symbolische heuristische Suche und zeigen, dass die Suchleistung nicht immer von der Verwendung einer Heuristik profitiert und dass die Suchleistung selbst unter den bestmöglichen Umständen, nämlich der perfekten Heuristik, exponentiell abnehmen kann. Dies unterstreicht, dass die symbolische blinde Suche heutzutage die dominierende symbolische Suchstrategie ist, ähnlich gut wie anderen modernen kostenoptimalen Planungsstrategien.

Basierend auf dieser Beobachtung und dem Mangel an guten Heuristiken für Planungsformalisten mit ausdrucksstarken Erweiterungen, erweist sich die symbolische Suche als ein starker Ansatz. Wir erweitern die symbolische Suche, um jeden der Formalismen einzeln und in Kombination zu unterstützen, was zu optimalen, korrekten und vollständigen Planungsalgorithmen führt, die sich im Vergleich zu anderen Ansätzen empirisch besser verhalten.

Acknowledgments

This thesis is the final product, which in itself required a lot of effort, but the years during the doctoral process laid the foundation for everything. During this time, I have had the pleasure of meeting, learning from, researching with, and being supported by many great people, for which I would like to express my gratitude.

First of all, I would like to thank Bernhard Nebel, who not only gave me the opportunity to do my own research that interested me, but at the same time kept inspiring me with his expertise and dedication to topics like complexity theory. The Foundations of Artificial Intelligence led by Bernhard was a great place to work with many wonderful people such as Florian Geißer, Gregor Behnke, Grigorios Mouratidis, Johannes Aldinger, Petra Geiger, Robert Mattmüller, Rolf-David Bergdoll, Stefan Staeglich, Thorsten Engesser, Tim Schulte and Uli Jakob, all of whom I would like to thank for great discussions and support.

I would especially like to thank Robert Mattmüller, who has supported and guided me since my first academic steps. There are few people who can inspire young people with great commitment and enthusiasm for science. Among them, there are even fewer that allow young scientists to be seen as independent researchers by not putting themselves in the center of the action. Robert is one of those few people without whom this work would not have been possible. Thank you very much.


I would like to thank all my colleagues with whom I have had the pleasure of writing, revising, and publishing a paper during my doctoral process: Álvaro Torralba, André Biedenkapp, Bernhard Nebel, David Borukhson, Dominik Drexler, Florian Geißer, Frank Hutter, Gregor Behnke, Jendrik Seipp, Marius Lindauer, Michael Katz, Robert Mattmüller, Sumitra Corraya and Thomas Keller. A special thanks to the planning community that has welcomed me since I first attended ICAPS 2018. In particular, to Álvaro Torralba for many interesting discussions, and to Michael Katz for the opportunity to work with him and visit the IBM T. J. Watson Research Center in New York. Both treated me like a

colleague from the first meeting and shared their expertise with me.

Being a doctoral student is a roller coaster ride that is only possible with a well-balanced social life. I would like to take this opportunity to thank all my friends with whom I have shared many great experiences in my gymnastics career. Thanks to my study friends André Biedenkapp, Lukas Gemein and Rick Gelhausen, all of whom will (hopefully) soon receive their PhDs. Special thanks to André Biedenkapp, with whom I have shared many personally great times and scientific successes but also some setbacks. The many overtime hours with you became a pleasure thanks to your positive attitude and enthusiasm for science.

I would like to thank my family. My parents Thomas and Olga, who have always supported me and taught me from an early age that hard work will pay off in the long run. They together with Iva, Mario and their slowly but steadily growing family were the support that many deserve but few receive. I am grateful that I can consider you not only family but also friends.

Last but not least, I would like to thank Lena (and Balu), who joined my journey a little later but are now indispensable. It is hard to describe how much I enjoyed the last time, although it should actually be very exhausting to write such a thesis. Lena, your incredibly loving, understanding and motivating nature towards me and my work has helped me to accomplish this thesis.



Contents

1	Introduction	1
1.1	Outline	3
1.2	Published Work	4
1.3	Awards	6
2	Background to Classical Planning	7
2.1	Formalism	7
2.2	Complexity and Compilations	10
2.3	Symbolic Search	11
2.4	Decision Diagrams	12
3	Symbolic Heuristic Search	15
3.1	Theoretical Results	17
3.2	Empirical Results	19
3.3	Discussion	21
4	Axioms and Derived Variables	23
4.1	Formalism	25
4.2	Complexity and Compilability	26
4.3	Symbolic Search	27
5	State-Dependent Action Costs	31
5.1	Formalism	33
5.2	Complexity and Compilability	34
5.3	Symbolic Search	36
6	Oversubscription Planning	39
6.1	Formalism	41
6.2	Complexity and Compilability	42
6.3	Symbolic Search	44

7 Top-k Planning	47
7.1 Formalism	48
7.2 Complexity and Compilability	49
7.3 Symbolic Search	50
8 Discussion	55
8.1 Combination of Extensions	55
8.2 Future Work	58
9 Conclusion	61
List of Figures	63
List of Tables	65
Bibliography	67

Introduction

If you fail to plan, you are planning to fail.

— Benjamin Franklin

Automated planning is the science of designing and engineering machines, in particular computer programs, that can automatically derive behaviors to achieve goals. The generation of such strategies, i.e., thinking before acting, is understood as an intelligent behavior and is one of the oldest areas in the field of Artificial Intelligence. Nowadays, a classical planning problem can be informally described as the problem of finding a course of actions that allows an intelligent agent to move from any situation it finds itself in to one that satisfies its goals. Since planning is not limited to a specific application domain, it was originally labeled as general problem solving (Helmert 2008; Newell and Simon 1963) and can be used for different types of reasoning problems, such as elevator control (Koehler and Schuster 2000), greenhouse logistics (Helmert and Lasinger 2010), natural-language generation (Koller and Hoffmann 2010), robot control (Karpas and Magazzeni 2020; Nilsson 1984; Speck, Dornhege, and Burgard 2017), risk management (Katz, Srinivas, et al. 2021; Sohrabi, Riabov, Katz, et al. 2018), and many more.

In classical planning, some properties of a planning problem at hand require an expressive extension of the standard classical planning formalism to capture them in a concise way. In this thesis, we consider four different extensions of classical planning (Figure 1.1), all of which capture and extend different aspects of classical planning while retaining the core of the formalism: Single-agent planning problems in a fully observable, static, discrete, deterministic, and fully known environment (Russell and Norvig 2003). Unfortunately, most classical planners do not support any of these expressive extensions because they are usually based on heuristic search and it appears very challenging to design informative and fast to compute heuristics (goal-distance estimators) that consider additional problem properties. This is especially true for cost-

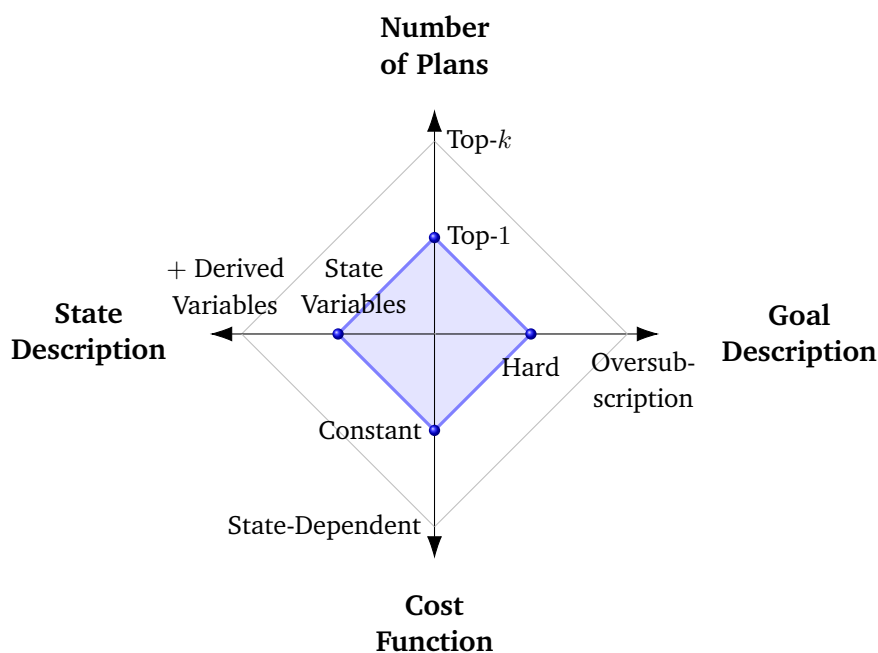


Figure 1.1: Overview of extensions for classical planning, where the blue color denotes the planning formalism supported by symbolic search approaches from the literature.

optimal planners, which additionally require that a heuristic is admissible, i.e., that the heuristic never overestimates the cost of reaching a goal state. However, it is well known that such extensions are critical for an elegant and compact modeling of many real-world problems. For example, extending the state description to include derived variables allows aspects of a planning problem that are not directly affected by the actions but are derived from the values of other variables to be modeled concisely using a set of logical axioms (Thiébaux, Hoffmann, and Nebel 2005). While in probabilistic planning a concise encoding of action costs or rewards in form of Markov decision processes has long been standard, in classical planning it is common to consider a potentially exponentially larger representation with constant action costs (Geißer 2018). In many real-world applications, the description of goals is oversubscribed, i.e., there are a large number of desirable, often competing goals of varying value, and a system may not be able to achieve all of them with the available resources (Smith 2004). Unfortunately, conventional classical planning cannot handle such a scenario, because the goal is treated as a hard constraint that can either be accomplished or not. Finally, in some cases, a single plan may be sufficient, but in practice it is often better to have many good alternatives to choose from in order to be more flexible (Nguyen et al. 2012).

In this thesis, we investigate and discuss the computational complexity and compilability of classical planning incorporating the four expressive extensions mentioned earlier. This analysis shows and highlights that for many real-world problems, it is desirable and even necessary to support these features natively. We propose for the first time to use symbolic search for cost-optimal planning for those four expressive extensions (see Figure 1.1). Symbolic search provides a good basis for native support of these planning formalisms for several reasons. A theoretical and empirical analysis of symbolic heuristic search in the form of BDDA^{*} shows that the use of a heuristic in symbolic search does not always improve search performance. This observation strengthens the fact that nowadays symbolic blind search, i.e., without heuristics, is the dominant search strategy of symbolic search, on par with explicit heuristic search in cost-optimal planning. Thus, since symbolic search does not necessarily need heuristics to be efficient, the search efficiency does not suffer from the lack of efficient and informative heuristics for the proposed planning scenarios. Based on these observations, we enhance symbolic search to obtain optimal, sound, and complete algorithms for planning with the expressive extensions under consideration. Our empirical evaluations show that the presented symbolic search approaches perform favorably in all these planning settings compared with other state-of-the-art approaches. Finally, we show that the proposed symbolic search approach is able to support planning tasks that use and require all expressive extensions at once.

1.1 Outline

This thesis is structured as follows. In Chapter 2, we introduce the relevant background for this thesis concerning classical planning. In particular, we introduce the concepts of complexity, compilability and expressive power and present symbolic search for classical planning together with different types of decision diagrams. In Chapter 3, we investigate the question, why heuristics do not seem to pay off in symbolic search. Specifically, we theoretically and empirically analyze the search behavior of symbolic heuristic search in form of BDDA^{*} (Speck, Geißer, and Mattmüller 2020). In Chapter 4, we summarize computational complexity and compilability results for planning with axioms from the literature. We introduce three ways to extend symbolic search algorithms to support axioms natively and present an empirical evaluation (Speck, Geißer, Mattmüller, and Torralba 2019). In Chapter 5, we present computational complexity and compilability results for planning with state-dependent action costs (Speck, Borukhson, et al. 2021). Then, symbolic search algorithms for planning with state-dependent action costs are presented and an empirical evaluation is conducted (Speck, Geißer, and Mattmüller 2018a). In Chapter 6, we discuss computational complexity and compilability results for oversubscription planning. A symbolic search approach is presented and explained for

oversubscription planning and an empirical evaluation is conducted (Speck and Katz 2021). In Chapter 7, we present computational complexity and compilability results for top- k planning and introduce symbolic search for it, which we analyze theoretically and empirically (Speck, Mattmüller, and Nebel 2020). In Chapter 8, we discuss the combination of the previously analyzed extensions for classical planning and the use of symbolic search for such a setting. Finally, possible future work related to this thesis is discussed. Chapter 9 concludes and summarizes this thesis.

1.2 Published Work

The core results presented in this work have been published at leading AI and automated planning conferences. In particular, the following publications form the backbone of this work. At the beginning of each chapter, we indicate which publications form the basis for the content of the chapter.

- **David Speck**, David Borukhson, Robert Mattmüller, and Bernhard Nebel (2021). “On the Compilability and Expressive Power of State-Dependent Action Costs”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 358–366. **(Best Student Paper Runner-Up Award)**
- **David Speck** and Michael Katz (2021). “Symbolic Search for Oversubscription Planning”. In: *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*. Ed. by Kevin Leyton-Brown and Mausam. AAAI Press, pp. 11972–11980.
- **David Speck**, Florian Geißer, and Robert Mattmüller (2020). “When Perfect Is Not Good Enough: On the Search Behaviour of Symbolic Heuristic Search”. In: *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)*. Ed. by J. Christopher Beck, Erez Karpas, and Shirin Sohrabi. AAAI Press, pp. 263–271.
- **David Speck**, Robert Mattmüller, and Bernhard Nebel (2020). “Symbolic Top- k Planning”. In: *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*. Ed. by Vincent Conitzer and Fei Sha. AAAI Press, pp. 9967–9974.
- **David Speck**, Florian Geißer, Robert Mattmüller, and Álvaro Torralba (2019). “Symbolic Planning with Axioms”. In: *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*. Ed. by Nir Lipovetzky, Eva Onaindia, and David E. Smith. AAAI Press, pp. 464–472.

- **David Speck**, Florian Geißer, and Robert Mattmüller (2018a). “Symbolic Planning with Edge-Valued Multi-Valued Decision Diagrams”. In: *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*. Ed. by Mathijs de Weerd, Sven Koenig, Gabriele Röger, and Matthijs Spaan. AAAI Press, pp. 250–258. **(Partly based on ideas from my master thesis)**

The following publications also resulted from my doctoral research, but are not a central part of this thesis. However, some of the ideas and concepts presented in this thesis were used in these research papers.

- **David Speck**, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (2021b). “Learning Heuristic Selection with Dynamic Algorithm Configuration”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 597–605.
 - **David Speck**, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (2020b). “Learning Heuristic Selection with Dynamic Algorithm Configuration”. In: *ICAPS 2020 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL)*, pp. 61–69. **(superseded)**
- Dominik Drexler, Jendrik Seipp, and **David Speck** (2021). “Subset-Saturated Transition Cost Partitioning”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 131–139.
 - Dominik Drexler, **David Speck**, and Robert Mattmüller (2020). “Subset-Saturated Transition Cost Partitioning for Optimal Classical Planning”. In: *ICAPS 2020 Workshop on Heuristics and Search for Domain-independent Planning (HSDIP)*, pp. 23–31. **(superseded)**
- Gregor Behnke and **David Speck** (2021). “Symbolic Search for Total-Order HTN Planning”. In: *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*. Ed. by Kevin Leyton-Brown and Mausam. AAAI Press, pp. 11744–11754.
- Florian Geißer, **David Speck**, and Thomas Keller (2020). “Trial-Based Heuristic Tree Search for MDPs with Factored Action Spaces”. In: *Proceedings of the 13th Annual Symposium on Combinatorial Search (SoCS 2020)*. Ed. by Daniel Harabor and Mauro Vallati. AAAI Press, pp. 38–47.

- Sumitra Corraya, Florian Geißer, **David Speck**, and Robert Mattmüller (2019). “An Empirical Study of the Usefulness of State-Dependent Action Costs in Planning”. In: *Proceedings of the 42nd Annual German Conference on Artificial Intelligence (KI 2019)*. Ed. by Christoph Benzmüller and Heiner Stuckenschmidt. Lecture Notes in Artificial Intelligence. Springer, pp. 123–130.
- Florian Geißer, **David Speck**, and Thomas Keller (2019). “An Analysis of the Probabilistic Track of the IPC 2018”. In: *ICAPS 2019 Workshop on the International Planning Competition (WIPC)*, pp. 27–35.
- **David Speck**, Florian Geißer, and Robert Mattmüller (2018b). “SYMPLE: Symbolic Planning Based on EVMDDs”. In: *Ninth International Planning Competition (IPC-9): Planner Abstracts*, pp. 91–94.
- Florian Geißer and **David Speck** (2018). “Prost-DD – Utilizing Symbolic Classical Planning in THTS”. in: *Sixth International Probabilistic Planning Competition (IPPC-6): Planner Abstracts*, pp. 13–16. (**Winner of the Discrete Markov Decision Process Track**)

Finally, the following patents were filed based on work conducted during my doctoral process.

- **David Speck**, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (Apr. 28, 2021a). “Device and Method for Planning and Operation of a Technical System”. U.S. pat. req. 17/242,790. Robert Bosch GmbH. Filed.
- **David Speck**, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (June 5, 2020a). “Device and Method for Planning and Operation of a Technical System”. European pat. req. EP 20178576.3 – 1203. Robert Bosch GmbH. Filed.

1.3 Awards

Together with my colleagues, I received the following awards for my work during my doctoral process.

- **Best Student Paper Runner-Up Award** for the paper “On the Compilability and Expressive Power of State-Dependent Action Costs” (**Speck**, Borukhson, et al. 2021) at ICAPS 2021
- **Winner of the Discrete Markov Decision Process Track** for the planning system “Prost-DD” (Geißer and **Speck** 2018) at the Sixth International Probabilistic Planning Competition at ICAPS 2018

Background to Classical Planning

[Binary Decision Diagrams are] one of the only really fundamental data structures that came out in the last twenty-five years.

— Donald Knuth (2008)

In AI planning, the objective is to automatically find a solution (a plan) to a given problem (a planning task). In this chapter, we formally define classical planning and discuss the concepts of complexity, compilability and expressive power in this context. Finally, symbolic search for classical planning is presented and different types of decision diagrams are introduced.

2.1 Formalism

In this thesis, we consider classical planning domains and tasks that are characterized by the SAS^+ formalism (Bäckström and Nebel 1995).

Definition 1 (Planning Domain). A *planning domain* is a tuple $\Xi = \langle \mathcal{V}, \mathcal{O}, \mathcal{C} \rangle$. \mathcal{V} is a finite set of state variables, each associated with a finite domain $D_v = \{0, \dots, |D_v| - 1\}$. A *fact* is a pair (v, d) , where $v \in \mathcal{V}$ and $d \in D_v$. For binary variables we also write v for $(v, 1)$ and $\neg v$ for $(v, 0)$. A *partial state* or *partial variable assignment* s over \mathcal{V} is a function on some subset of \mathcal{V} such that $s(v) \in D_v$, wherever $s(v)$ is defined. If s assigns a value to each variable $v \in \mathcal{V}$, s is called a *state*. With \mathcal{S} we refer to the set of all possible states defined over \mathcal{V} . \mathcal{O} is a finite set of *operators/actions*, where each operator is a pair $o = \langle \text{pre}_o, \text{eff}_o \rangle$ of partial variable assignments, called *preconditions* and *effects*. An operator $o \in \mathcal{O}$ is applicable in a state s iff pre_o is satisfied in s , i.e., $\text{pre}_o \subseteq s$. Applying operator o to state s yields the state $s[o] = s'$, where $s'(v) = \text{eff}_o(v)$ for all variables $v \in \mathcal{V}$ for which eff_o is defined and $s'(v) = s(v)$ otherwise. Finally, $\mathcal{C} : \mathcal{O} \rightarrow \mathbb{N}_0$ is the *cost function* of Ξ . The size

of Ξ is $\|\Xi\| = \sum_{v \in \mathcal{V}} |D_v| + \sum_{o \in \mathcal{O}} (|pre_o| + |eff_o|) + \|\mathcal{C}\|$, where $\|\mathcal{C}\|$ is bounded by $|\mathcal{O}| \cdot \lceil \log_2 N \rceil$, where N is the highest action cost value.

Planning domains together with initial states and goal descriptions form planning tasks.

Definition 2 (Planning Task). A *planning task* is a tuple $\Pi = \langle \Xi, \mathcal{I}, \mathcal{G}, \mathcal{B} \rangle$ consisting of a *planning domain*, an *initial state* $\mathcal{I} \in \mathcal{S}$, a partial variable assignment \mathcal{G} (*goal condition*), which defines all possible goal states $S_\star \subseteq \mathcal{S}$, and a *cost bound* $\mathcal{B} \in \mathbb{N}_0 \cup \{\infty\}$. For simplicity, we also write $\langle \mathcal{V}, \mathcal{O}, \mathcal{C}, \mathcal{I}, \mathcal{G} \rangle$ for a planning task with domain $\Xi = \langle \mathcal{V}, \mathcal{O}, \mathcal{C} \rangle$ and $\mathcal{B} = \infty$. The *size* of Π is $\|\Pi\| = \|\Xi\| + |\mathcal{G}| + |\mathcal{I}| + \|\mathcal{B}\|$, where $\|\mathcal{B}\| = \lceil \log_2 \mathcal{B} \rceil$ is the binary encoding size of the cost bound.

A planning task is often defined without a cost bound \mathcal{B} , which is equivalent to setting the cost bound to infinity (or a sufficient upper bound). However, we use planning tasks with cost bounds to perform complexity and compilability analyses in this dissertation.

The objective of classical planning is to determine plans, which are sequences of applicable actions leading from the initial state to a goal state, respecting the given cost bound \mathcal{B} .

Definition 3 (Plan). A *plan* $\pi = \langle o_0, \dots, o_{n-1} \rangle$ for planning task Π is a sequence of applicable operators that generates a sequence of states s_0, \dots, s_n , where $s_0 = \mathcal{I}$, $s_n \in S_\star$ is a goal state, $s_{i+1} = s_i[o_i]$ for all $i = 0, \dots, n-1$, and the *cost* of π is less than or equal to the cost bound \mathcal{B} , i.e., $cost(\pi) = \sum_{i=0}^{n-1} \mathcal{C}(o_i) \leq \mathcal{B}$. Plan π is *optimal* if there is no cheaper plan. We call $\|\pi\| = n$ the *length* or *size* of π . With P_Π we refer to the (possibly infinite) *set of all plans* for a planning task Π .

The search for an optimal plan, i.e., a plan with the lowest cost, is called cost-optimal planning, or optimal planning for short, and is the focus of this thesis. Example 1 describes a classical planning domain and task of a Mars rover, which we will use in this work to motivate and explain expressive features that extend the classical planning formalism.

Example 1. Consider a Mars rover like Perseverance¹ equipped with a drone like Ingenuity¹, which are designed to perform autonomous tasks. Such a scenario is illustrated in Figure 2.1. The dynamics of this example, i.e., the domain, are as follows.

The rover can navigate between adjacent cells if they are free (impassable cells are highlighted in red). Navigating the rover between cells has no cost, i.e., a cost of 0. The drone can launch and land on the rover if both are in the same location. Launching and landing each has a cost of 5. In addition, the

¹<https://mars.nasa.gov/mars2020/mission/overview/> (Accessed: 2021-10-12)




Figure 2.1: Visualization of the Mars rover planning task used as a running example. The original image is from NASA/JPL-Caltech/University of Arizona² and shows the Jezero crater, where the green dot indicates the actual landing site of the Perseverance rover. Rover, drone, grid lines, red coloring (impassable cells for the rover), camera (goal: take pictures), and arrow (goal: cell to which the rover should travel with the drone) are added.

drone can take an image at its current position at a cost of 2 and fly between any two cells, ignoring the impassability of cells to the rover. The cost of flying from one cell $from = (x, y)$ to another cell $to = (x', y')$ is the Manhattan distance between these two cells, i.e., $|x - x'| + |y - y'|$.

In this particular planning task, the rover along with the drone is initially located at (7,3), the actual landing site of Perseverance. The goal is to take pictures at (6,1) and (10,1) and travel with the drone equipped to (0,5), a location known as “Three Forks” from which new missions can be commenced. The cost bound for this task is infinite.

An optimal plan for this task is $\pi^r = \langle \text{navigate-7-2}, \text{navigate-7-1}, \text{launch-7-1}, \text{fly-7-1-to-6-1}, \text{take-img-6-1}, \text{fly-6-1-to-10-1}, \text{take-img-10-1}, \text{fly-10-1-to-7-1}, \text{land-7-1}, \text{navigate-7-2}, \text{navigate-7-3}, \dots, \text{navigate-0-5} \rangle$ with a cost of $\text{cost}(\pi^r) = 5 + 5 + 2 \cdot 2 + 8 = 24$, since navigating the rover costs 0, the drone launches and lands once with a cost of 5 each, takes images twice with a cost of 2 each, and flies a total distance of 8 cells.

²<https://mars.nasa.gov/resources/25621/perseverances-landing-spot-in-jezero-crater/> (Accessed: 2021-10-12)

2.2 Complexity and Compilations

A natural question that arises is how hard planning is in general. More precisely, *the bounded plan existence problem* of planning is the problem of deciding for a given planning task Π whether there exists a plan. It was shown that this problem is PSPACE-complete for STRIPS (Bylander 1997) and SAS⁺ planning tasks (Bäckström and Nebel 1995).

Theorem 1 (Bylander 1994). *Bounded plan existence of planning is PSPACE-complete.* \square

Interestingly, if two planning formalisms belong to the same complexity class, it does not directly mean that both have the same expressive power. Expressive power is a measure of how concisely planning domains and plans can be expressed in a given formalism with compilation schemes (Nebel 2000). In the further course of this work, we will use compilation schemes to show whether and under which conditions certain model extensions such as derived variables or state-dependent action costs can be compiled away, i.e., expressed in the original formalism as a classical planning task.

The following definition formalizes compilation schemes, which translate from one planning formalism to another while preserving plan existence and polynomially preserving task sizes.

Definition 4 (Compilation Scheme). A compilation scheme or, in short, a compilation is a tuple of functions $\mathbf{f} = \langle f_\xi, f_\iota, f_g \rangle$ on planning domains that induces a function on planning tasks as follows:

$$F(\Pi) = \langle f_\xi(\Xi), \mathcal{I} \cup f_\iota(\Xi), \mathcal{G} \cup f_g(\Xi), \mathcal{B} \rangle,$$

where $\Pi = \langle \Xi, \mathcal{I}, \mathcal{G}, \mathcal{B} \rangle$, satisfying the following conditions:

1. there exists a plan for $F(\Pi)$ iff there exists a plan for Π , and
2. the size of the results of f_ξ , f_ι , and f_g is polynomial in the size of the arguments.

Note that our definition of a compilation is “cost-sensitive”. By not allowing the cost bound to change, we ensure that for every plan in the original task, there must be a plan with cost at most \mathcal{B} in the target task, and vice versa. Besides preserving planning task sizes polynomially, another desirable property is preservation of plan lengths. This is captured by the following definition.

Definition 5 (Compilations Preserving Plan Length). A compilation scheme \mathbf{f} is said to *preserve plan length exactly* (modulo an additive constant) if for every plan π solving an instance Π , there exists a plan π' solving $F(\Pi)$ with $\|\pi'\| \leq \|\pi\| + k$ for some constant $k \in \mathbb{N}_0$. It is said to *preserve plan length linearly* if $\|\pi'\| \leq c \cdot \|\pi\| + k$ for constants $c \in \mathbb{N}_0$ and $k \in \mathbb{N}_0$, and to *preserve plan length polynomially* if $\|\pi'\| \leq p(\|\pi\|, \|\Pi\|)$ for some polynomial p .

Intuitively, compilability while exactly preserving the plan length shows that the planning formalism we use as the target formalism is at least as expressive as the source formalism. If compilation requires polynomial or even exponential plan growth, this can lead to an infeasible challenge for planning algorithms that indicates an increase in expressive power (Nebel 2000; Thiébaux, Hoffmann, and Nebel 2005).

2.3 Symbolic Search

Symbolic search is a state space exploration technique that has its origin in the field of Model Checking (McMillan 1993). Symbolic search algorithms resemble their explicit counterparts, but expand and generate whole sets of states in contrast to individual states. In symbolic search, a set of states $S \subseteq \mathcal{S}$ is represented by its *characteristic function* χ_S , which is a Boolean function $\chi_S : \mathcal{S} \rightarrow \{0, 1\}$ that represents whether a given state belongs to S or not. More precisely, states contained in S are mapped to 1 and all others to 0, i.e., $\chi_S(s) = 1$ if $s \in S$ and $\chi_S(s) = 0$ otherwise. Similarly, operators \mathcal{O} can be represented as so-called *transition relations (TRs)*, which are sets of state pairs, namely predecessor and successor states. The characteristic function of a transition relation T_O representing a set of operators $O \subseteq \mathcal{O}$ is a function $\chi_{T_O} : \mathcal{S} \times \mathcal{S} \rightarrow \{0, 1\}$ that maps all pairs of states (s, s') to true iff successor $s' \in \mathcal{S}$ is reachable from predecessor $s \in \mathcal{S}$ by applying an operator $o \in O$. Given a set of states S and a TR T_O , the *image (preimage)* operator computes the set of successor (predecessor) states S' of S through T_O . Note that a single TR can in general represent multiple operators with the same cost (Torralba, Alcázar, et al. 2017; Torralba, Edelkamp, and Kissmann 2013).

Symbolic (blind) search describes a symbolic version of uniform cost search, also known as Dijkstra’s algorithm (Dijkstra 1959), which can be performed in different search directions. *Symbolic forward (blind) search* (progression) starts from the representation of the initial state χ_I , and iteratively computes the image until a set of states S is found whose intersection with the goal χ_G is non-empty, i.e., $\chi_S \wedge \chi_G \neq \perp$. The *open* and *closed list* are represented as lists of state sets partitioned into subsets with identical g -values, where the g -value describes the cost required to reach these states. During the search, the closed list is used to track and prune states that have already been expanded. *Symbolic backward (blind) search* (regression) can be realized by starting with the goal states, applying the preimage operation until the initial state is found. In *symbolic bidirectional (blind) search*, both forward and backward symbolic search are performed simultaneously, maintaining two symbolic searches with separate open and closed lists. A search step consists either of a backward or a forward search step (and modifies the respective open and closed lists). If a state of the current search direction is expanded, which is already contained in the closed list of the search in the opposite direction, a goal path is found. In

general, all strategies, which switch iteratively between both search directions, guarantee optimality if the termination criterion is chosen accordingly (Pohl 1969).

Finally, a *plan reconstruction* (Torralba 2015) is performed to obtain the final plan. In explicit search, each (search) node keeps track of its parent node, making it easy to construct a plan when a goal state is found. In symbolic search, however, the parents are not directly known, but all parents are stored in the closed list with their reachability costs. Therefore, it is possible to perform a greedy search, which opposes the actual search direction with the perfect heuristic (cost of shortest path) obtained by the corresponding closed list. In forward (backward) symbolic search the plan is constructed by a greedy backward (forward) search starting with a found goal state (the initial state). The plan reconstruction procedure iterates over all operators (descending cost) and selects an explicit predecessor (successor) contained in the closed list. The latter process is repeated until the initial state (a goal state) is reached. Note that a greedy search in combination with the perfect heuristic leads the search directly from a starting state to a target state, making the runtime of the plan reconstruction negligible with respect to the actual search. For bidirectional search, a greedy best-first search is performed twice, both opposing the actual search direction. More specifically, both plan reconstructions are initialized with the meeting point and one search is a regression to the initial state, while the other search is a progression to the goal states.

2.4 Decision Diagrams

In symbolic search the most prominent way to represent (characteristic) functions are decision diagrams such as Binary Decision Diagrams (BDDs) (Bryant 1985), Algebraic Decision Diagrams (ADDs) (Bahar et al. 1997) or Edge-Valued Binary Decision Diagrams (EVBDDs) (Lai, Pedram, and Vrudhula 1996) which are data structures that offer a compromise between conciseness of representation and efficiency of manipulation (Drechsler and Becker 1998a). The main idea is to break down a function f into subfunctions, so that f can be reassembled from them.

Definition 6 (Binary Decision Diagram). A *Binary Decision Diagram (BDD)* B_S is a directed acyclic graph with a single root node and two terminal nodes: the 0-sink and the 1-sink. Each inner node corresponds to a binary³ variable $v \in \mathcal{V}$ and has two successors, where the *low edge* represents that variable v is false, while the *high edge* represents that variable v is true. By traversing the

³Note that each finite domain variable $v \in \mathcal{V}$ can be represented by $\lceil \log_2 |D_v| \rceil$ binary variables. Although there are generalizations for each type of decision diagram to support multi-valued variables directly, this dissertation mainly considers decision diagrams that support only binary variables, unless explicitly stated otherwise.

BDD according to a given assignment, the represented Boolean function χ_S can be evaluated.

ADDs and EVBDDs have been successfully used to represent numerical functions $f : \mathcal{S} \rightarrow \mathbb{Q} \cup \{\infty\}$ (Hansen, Zhou, and Feng 2002; Speck, Geißer, and Mattmüller 2018b; Torralba, Linares López, and Borrajo 2013) in symbolic planning. An *Algebraic Decision Diagrams (ADD)* A_f is similar to a BDD, but has an arbitrary number of terminal nodes with different discrete values including real numbers. An *Edge-Valued Binary Decision Diagram (EVBDD)* E_f is a rooted directed acyclic graph with edge values, a dangling incoming edge to the root node and a single terminal node. The function f can be evaluated by traversing the graph according to the variable assignment and simultaneously adding up the edge weights. The resulting sum is the function value for the corresponding variable assignment. In practice, the generalization of EVBDDs, so-called *Edge-Valued Multi-valued Decision Diagrams (EVMDDs)* (Ciardo and Siminiceanu 2002), where variables can be multi-valued, is more common in planning than EVBDDs (Geißer, Keller, and Mattmüller 2015, 2016; Mattmüller et al. 2018).

Decision diagrams are typically considered in a reduced and ordered form (Becker and Molitor 2008) and can represent exponentially many states requiring only polynomial space. A decision diagram is called *ordered* if on all paths from the root to a sink variables appear in the same order. A decision diagram is called *reduced* if isomorphic subgraphs are merged and any node is eliminated whose two children are identical. For fixed variable orders, reduced and ordered decision diagrams are unique (Bahar et al. 1997; Bryant 1986; Lai, Pedram, and Vruthula 1996). Note that for EVMDDs the corresponding edge values must be taken into account. From now on we only talk about reduced and ordered decision diagrams and assume a fixed variable order.

Definition 7 (Decision Diagram Size). The *size* $|D|$ of a decision diagram D is the number of nodes of D .

The size of a decision diagram depends strongly on the variable order, so that a good order can lead to an exponentially more compact decision diagram (Edelkamp and Kissmann 2011). For some functions the size of the corresponding decision diagram is exponential, independent of the underlying variable order (Bryant 1986; Edelkamp and Kissmann 2011). Comparing the different types of decision diagrams, we can see that an EVBDD can be exponentially more compact than an ADD (Siminiceanu and Roux 2010) representing additively separable functions such as $f : \{0, 1\}^{n+1} \rightarrow \{0, \dots, 2^{n+1} - 1\}$ with $f(x_0, \dots, x_n) = \sum_{i=0}^n 2^i x_i$. Moreover, an ADD can be efficiently disassembled into multiple BDDs, one for each terminal node, in polynomial time and memory with respect to the ADD size by substituting terminal nodes (Torralba 2015). In practice, the main advantage of using BDDs over ADDs (and EVBDDs) is that decision diagram libraries such as CUDD (Somenzi 2015) use techniques

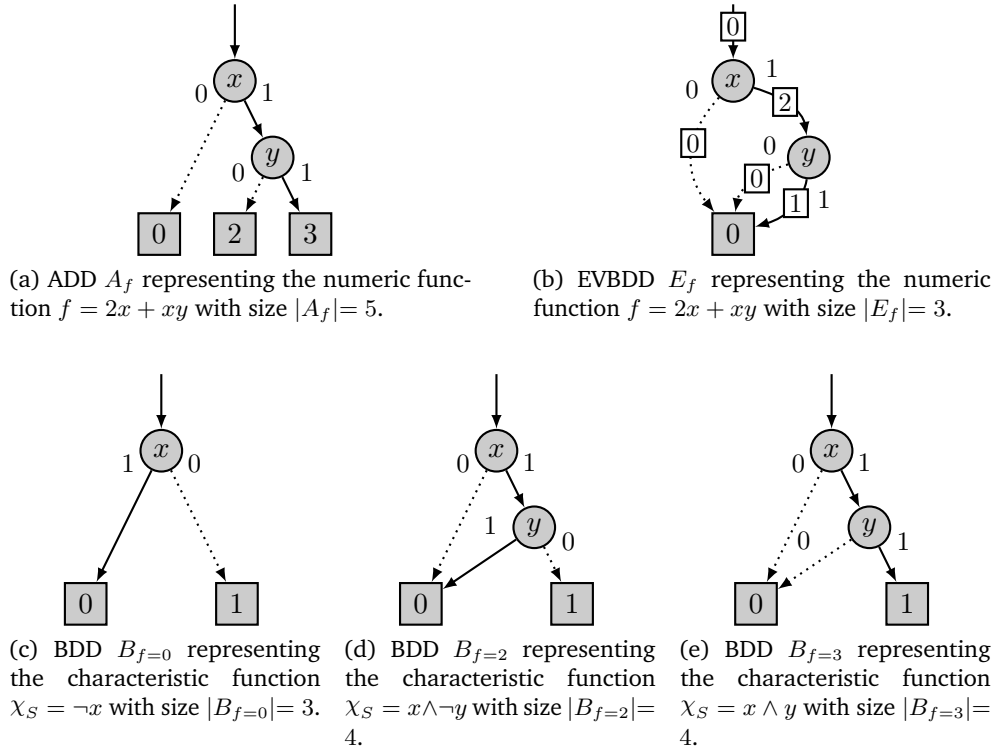


Figure 2.2: Visualization of different decision diagrams.

such as complement edges to store BDDs more compactly (Brace, Rudell, and Bryant 1990) and allow for more efficient operations (Burch et al. 1994).

Example 2. Figures 2.2a and 2.2b show the ADD A_f and EVBDD E_f representing the numeric function $f = 2x + xy$. The function f can also be represented as multiple BDDs by disassembling the ADD A_f into three different BDDs, one for each terminal node. Figures 2.2c to 2.2e depict the BDDs $B_{f=z}$ representing all states for which the evaluation of function $f = 2x + xy$ is z . The variable order for all decision diagrams is $x \succ y$, i.e., x appears before y on each path.

From now on, when we refer to state sets, characteristic functions and numerical functions, we assume that they are represented as one of the corresponding decision diagrams and all logical or numerical operations are realized with the efficient and appropriate decision diagram-based operations using the APPLY algorithm (Bahar et al. 1997; Bryant 1986; Lai, Pedram, and Vrudhula 1996). If it is important which type of decision diagram we have used to represent certain functions, we will specify it explicitly.

Symbolic Heuristic Search

While the main observations of this paper are both intuitive and pretty obvious, I still consider the work a significant contribution in focusing attention on a key and often overlooked difference between explicit and symbolic search [...].

— Reviewer #3 (2020)

Core Publication of this Chapter

- **David Speck**, Florian Geißer, and Robert Mattmüller (2020). “When Perfect Is Not Good Enough: On the Search Behaviour of Symbolic Heuristic Search”. In: *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)*. Ed. by J. Christopher Beck, Erez Karpas, and Shirin Sohrabi. AAAI Press, pp. 263–271.

Explicit search and symbolic search have been shown to be strong and competitive approaches for optimal classical planning. While explicit search usually occurs as variants of forward A^* search (Hart, Nilsson, and Raphael 1968), which are equipped with strong and efficient goal-distance heuristics, symbolic search is usually performed as (bidirectional) blind search, i.e., without heuristics. This naturally raises the question of why not combine the two techniques to obtain a planning algorithm that combines the state pruning of explicit heuristic search with the concise representation of symbolic search. This combination is referred to as symbolic heuristic search, which includes a variety of symbolic generalizations of A^* for different decision diagrams such as BDDA^{*} (Edelkamp and Reffel 1998), ADDA^{*} (Hansen, Zhou, and Feng 2002) and EVMDDA^{*} (Speck, Geißer, and Mattmüller 2018a).




Figure 3.1: Illustration of BDDA^{*}, where the cells represent state sets $S_{g,h}$ and the arrows denote successor state sets. The gray cells are expanded in the order indicated by the numbers (Torralba 2015; Torralba and Gnad 2017).

The underlying idea of symbolic A^{*} search is to precompute and represent a *heuristic function* h with decision diagrams, where $h : \mathcal{S} \rightarrow \mathbb{N}_0 \cup \{\infty\}$ such that $h(s_*) = 0$ for all $s_* \in \mathcal{S}_*$. Symbolic operations are then used to assign the corresponding f -values ($f = g + h$) to each state of a set of states S_g reachable with cost g . As usual with A^{*}, symbolic A^{*} expands the states in ascending order of the f -value and uses a tie-breaking rule in favor of states with smaller g -values (Kissmann and Edelkamp 2011; Torralba 2015). To guarantee optimal solutions, a *consistent heuristic* is assumed, i.e., its estimate is always less than or equal to the estimated distance from each direct successor to the goal, plus the cost to reach that successor, i.e., $h(s) \leq h(s[o]) + \mathcal{C}(o)$. Note that any consistent heuristic is also admissible, i.e., never overestimates the cost of reaching a goal state (Pearl 1984). Two important examples of consistent heuristics and at the same time the extreme cases are the perfect heuristic h_* , which maps any state s to the cost of the cheapest path from s to any goal state, and the blind heuristic h_{blind} , which maps any state s to 0.

In BDDA^{*}, a consistent heuristic function is precomputed and represented by multiple BDDs, one for each heuristic value h , where each BDD is used to represent the states S_h with heuristic value h . During search, states S_g reachable with cost g are partitioned according to their heuristic value by computing the intersections of S_g and S_h for each heuristic value h resulting in sets of states $S_{g,h}$. Example 3 exemplifies how BDDA^{*} works.

Example 3. We consider a planning task with unit costs. BDDA^{*} starts with the state set $S_{0,2} = \{\mathcal{I}\}$, which contains only the initial state with a g -value of 0 and an h -value of 2 (Figure 3.1a). The expansion of $S_{0,2}$ leads to sets of states

with g -value of 1 and different h -values. Next, the state set $S_{1,1}$ is expanded, then $S_{1,2}$, and so on until a goal state is contained in $S_{4,0}$ (Figure 3.1b).

While BDDA* utilizing a heuristic expands fewer states, symbolic blind search (BDDA* with h_{blind}) potentially expands fewer sets of states. Moreover, as Speck, Geißer, and Mattmüller (2020) have shown and we will examine in this chapter, the BDD representation of the state sets can deteriorate if the states are partitioned according to their h -values.

There are several heuristics that can be precomputed and represented using symbolic search and decision diagrams, leading to cutting-edge performance in explicit heuristic search (Edelkamp 2002; Franco, Lelis, and Barley 2018; Franco, Torralba, et al. 2017; Moraru et al. 2019). Thus, all the ingredients are present to allow a symbolic planner utilizing heuristics, as explicit planners do. However, Jensen, Veloso, and Bryant (2008) identified the partitioning of state sets according to their heuristic values as a bottleneck, because multiple arithmetic operations have to be performed during search. This leads to different extension of BDDA* to overcome this bottleneck such as Lazy BDDA* (Torralba 2015), which delays the heuristic evaluation as long as possible, or SetA* (Jensen, Veloso, and Bryant 2008), which encodes the heuristic values as preconditions of actions resulting in multiple actions with costs according to the heuristic values. However, empirical evaluations show that all versions of BDDA* perform better than symbolic blind search in some domains but overall symbolic bidirectional search without any heuristic performs best (Torralba, Linares López, and Borrajo 2016).

In the remainder of this chapter, we describe and summarize the results of Speck, Geißer, and Mattmüller (2020), which theoretically and empirically evaluate the search behavior of BDDA*. On the theoretical side, this study reveals another fundamental problem of symbolic heuristics, namely that the use of a heuristic does not always improve the search performance of BDDA*, as it may affect the size of the representation. In general, even the *perfect heuristic* can exponentially deteriorate the search performance of symbolic A*. The empirical evaluation is consistent with these theoretical results. Finally, we conclude this chapter with a discussion of the implications of these findings.

3.1 Theoretical Results

Speck, Geißer, and Mattmüller (2020) show that good goal-distance estimations in the form of heuristics are not the appropriate quantity to improve the search performance of symbolic heuristic search. The reason for this is that, in contrast to explicit A*, where every consistent heuristic can only reduce the number of necessary node expansions (up to tie-breaking) and thus the search effort compared to blind search, in BDDA* no such guarantee exists. In BDDA*, the size of expanded BDDs, i.e., number of BDD nodes, representing expanded states determines the search effort and thus the runtime. As a BDD $B_{S'}$ can

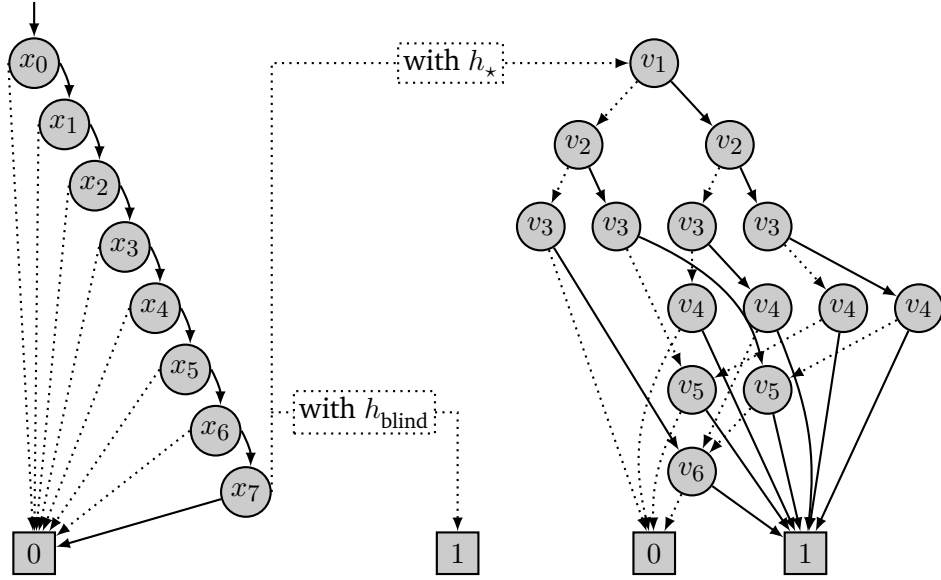


Figure 3.2: Visualization of two BDDs B_S (with h_{blind}) and $B_{S'}$ (with h_*) representing state sets S and S' , where $B_{S'}$ is exponentially larger in the number of variables than B_S , although $S' \subsetneq S$ (Speck, Geißer, and Mattmüller 2020).

be exponentially larger than a BDD B_S although the set of states S' is a strict subset of S , i.e., $S' \subsetneq S$, it is not always beneficial to represent and expand fewer states (Figure 3.2). In other words, in explicit search, where the most promising states (search fringe) are simply enumerated explicitly, reducing the number of states to expand directly improves search performance. However, in symbolic search, reducing the number of states to be expanded can lead to a larger representation size in form of BDDs and a fragmentation of the search fringe as the search progresses.

Similar to “must-expand” nodes in explicit search, Speck, Geißer, and Mattmüller (2020) introduce the notion of *expansion size* for BDDA^* as the cumulative size of BDDs that must always be expanded by BDDA^* before finding an optimal solution to measure search effort. Using this notion it is possible to prove that even under the best possible and unrealistic circumstances, namely the perfect heuristic, the search effort of BDDA^* can be exponentially larger than the search effort of symbolic search without heuristic, and vice versa (Theorem 2).

Theorem 2 (Speck, Geißer, and Mattmüller 2020). *Using the perfect heuristic h_* instead of the blind heuristic h_{blind} can decrease or increase the expansion size of BDDA^* exponentially in the size of the planning task for a given variable ordering.* \square

While the result that the search performance of BDDA^{*} can be exponentially improved when h_* is used instead of h_{blind} is less surprising, the opposite result is very surprising. This highlights the difference between explicit and symbolic search: the representation of the search fringe must be concise, and in symbolic search the decision diagram size is not directly related to the number of states represented. To prove that using h_* instead of h_{blind} can exponentially increase the expansion size and thus exponentially deteriorate search performance, Speck, Geißer, and Mattmüller (2020) constructed a family of planning tasks Π_n parameterized over the number of (relevant) variables. Solving tasks of this family of planning tasks Π_n with BDDA^{*} using the perfect heuristic prunes states from the expanded set of states (search fringe) so that the representation size increases exponentially. Figure 3.2 shows the core idea of the proof, with two BDDs representing the search fringe, i.e., the expanded state set, in the last expansion step of BDDA^{*} using h_* or h_{blind} when solving Π_3 . The key observation is that when h_{blind} is used, no states are pruned according to their values over variables v_i , resulting in a compact BDD representation. However, if h_* is used, states are pruned according to their values over the variables v_i and the function $(v_1 \wedge v_{n+1}) \vee \dots \vee (v_n \wedge v_{2n})$ has to be represented, which under certain variable orders requires a BDD with exponentially many nodes (Kissmann 2012).

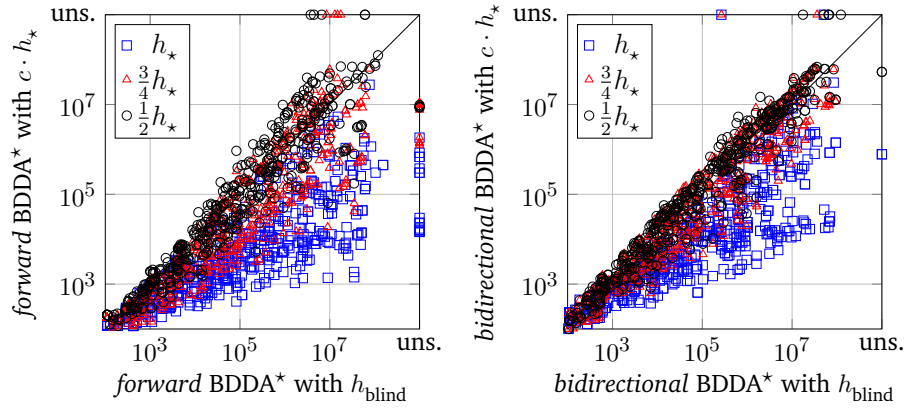
Finally, Speck, Geißer, and Mattmüller (2020) show that these theoretical results also hold for several other symbolic A^{*} variants such as Lazy BDDA^{*}, SetA^{*}, ADDA^{*}, and EVMDDA^{*}, as well as for their bidirectional extensions.

3.2 Empirical Results

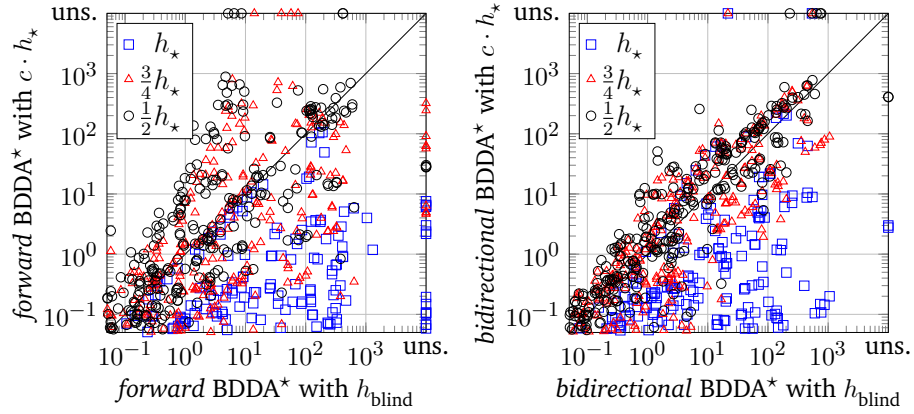
Speck, Geißer, and Mattmüller (2020) empirically investigated forward and bidirectional BDDA^{*} with precomputed fraction perfect heuristics on domains from the optimal track from the International Planning Competitions between 1998 and 2018. A heuristic $c \cdot h$ is called fraction perfect if it assigns to all states the values of the perfect heuristic multiplied by a constant $0 \leq c \leq 1$, where $0h_* = h_{\text{blind}}$ and $1h_* = h_*$ are important extreme cases. Note that BDDA^{*} with blind heuristic $0h_* = h_{\text{blind}}$ corresponds to symbolic blind search.

Figure 3.3a compares the search effort of BDDA^{*} with the blind heuristic and the search effort of BDDA^{*} with the fraction perfect heuristics. While the unrealistic case of the perfect heuristic almost always leads to a reduction in search effort in practice, when we consider more realistic cases, namely the fraction perfect heuristic, we can see that the search effort of BDDA^{*} can improve or deteriorate to the same extent.

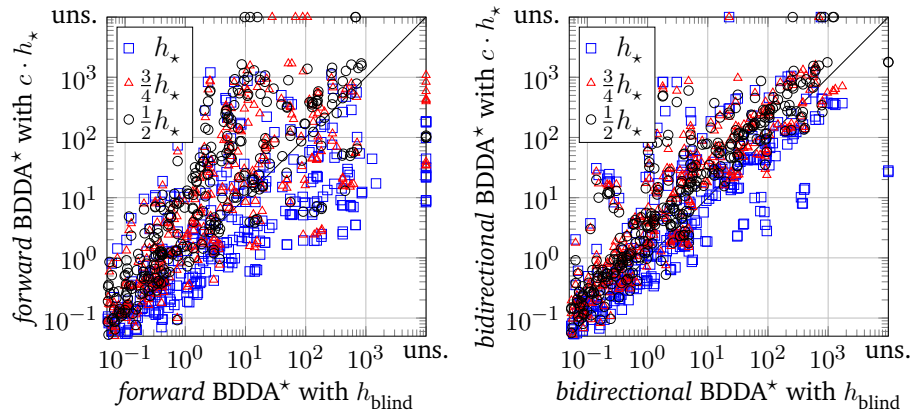
Figure 3.3b shows the expansion time of BDDA^{*}, i.e., the cumulative time required to generate all successors with the image operation. A correlation is observed between the search effort and the expansion time, which also empirically shows that the introduced notion of search effort is an adequate



(a) Search Effort (BDD nodes)



(b) Image Time (seconds)



(c) Overall Runtime (seconds)

Figure 3.3: Comparison of BDDA* with the blind heuristic and with given (fraction) perfect heuristics (Speck, Geißer, and Mattmüller 2020). With *uns.* we refer to instances that were not solved within the time and memory limits.

quantity to measure the performance of BDDA^{*}. This highlights the fundamental problem highlighted by Speck, Geißer, and Mattmüller (2020): using a heuristic does not always improve the search performance of BDDA^{*}.

Finally, Figure 3.3c shows the total running time of BDDA^{*} without the time for computing the corresponding heuristic. Comparing the expansion time with the total running time, we find that BDDA^{*} with fraction perfect heuristics has a higher time increase than BDDA^{*} with h_{blind} , which is due to the time-consuming partitioning of the state sets by heuristic values (Jensen, Veloso, and Bryant 2008).

Overall, these empirical results show that also in practice a heuristic can improve or deteriorate the search effort of BDDA^{*} to the same extent. These empirical results are consistent with the presented theoretical results. Furthermore, Speck, Geißer, and Mattmüller (2020) show that it appears to be domain dependent whether a heuristic helps BDDA^{*} as the structure of the reachable search space appears to play a central role.

3.3 Discussion

As the results of Speck, Geißer, and Mattmüller (2020) show, using goal-distance estimators, i.e., heuristics, to prune states in symbolic heuristic search does not always pay off. Indeed, even the perfect heuristic can exponentially deteriorate the search performance of BDDA^{*} and other symbolic A^{*} variants. Since the search performance of BDDA^{*} is not directly related to the number of expanded states, but to the size of the involved BDDs during the search, Speck, Geißer, and Mattmüller (2020) suggest using heuristic functions that can provide a nice structure of the involved decision diagrams or even give a size guarantee. One possible candidate are potential heuristics (Pommerening et al. 2015), which have recently been shown to be encodable as operator potentials within symbolic heuristic search, resulting in state-of-the-art performance (Fišer, Torralba, and Hoffmann 2021).

The fact that symbolic blind search does not use heuristics and yet can compete with explicit heuristic search in cost-optimal planning (Edelkamp, Kissmann, and Torralba 2015; Speck, Geißer, and Mattmüller 2020; Torralba, Alcázar, et al. 2017) has the advantage that it is not constrained by the limits of heuristics. Considering expressive extensions of classical planning, such as conditional effects, axioms, and many more, explicit heuristic search planners rarely support them. The reason is that it is very challenging to design admissible heuristics that are both informative and fast to compute when considering such extensions. Symbolic blind search, however, provides a good basis to support these extensions, since it does not rely on heuristics and performs an efficient exhaustive search. In the next chapters, we will take advantage of this and show how symbolic search can be adapted to support expressive extensions in cost-optimal planning.

Axioms and Derived Variables

The grand aim of all science [is] to cover the greatest number of empirical facts by logical deduction from the smallest possible number of hypotheses or axioms.

— Albert Einstein (1950)

Core Publication of this Chapter

- **David Speck**, Florian Geißer, Robert Mattmüller, and Álvaro Torralba (2019). “Symbolic Planning with Axioms”. In: *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*. Ed. by Nir Lipovetzky, Eva Onaindia, and David E. Smith. AAAI Press, pp. 464–472.

In classical planning, Boolean or finite domain variables are used to describe the states of the world, the preconditions and the effects of actions, and the goal description (Bäckström and Nebel 1995; Fikes and Nilsson 1971). For many real-world problems, however, complex action preconditions or goals are desirable or even necessary for a compact problem description (Thiébaux, Hoffmann, and Nebel 2005). Axioms allow to model such complex preconditions and goals compactly by introducing a set of derived variables whose values are not directly influenced by the actions but are derived from the values of other variables using a set of logical axioms.

Although axioms are an essential feature (Thiébaux, Hoffmann, and Nebel 2005) of PDDL (Hoffmann and Edelkamp 2005; McDermott et al. 1998), the common language for modeling planning tasks, modern planning techniques rarely support axioms, especially in cost-optimal planning. Most admissible heuristics commonly used in A* search, one of the most prominent approaches

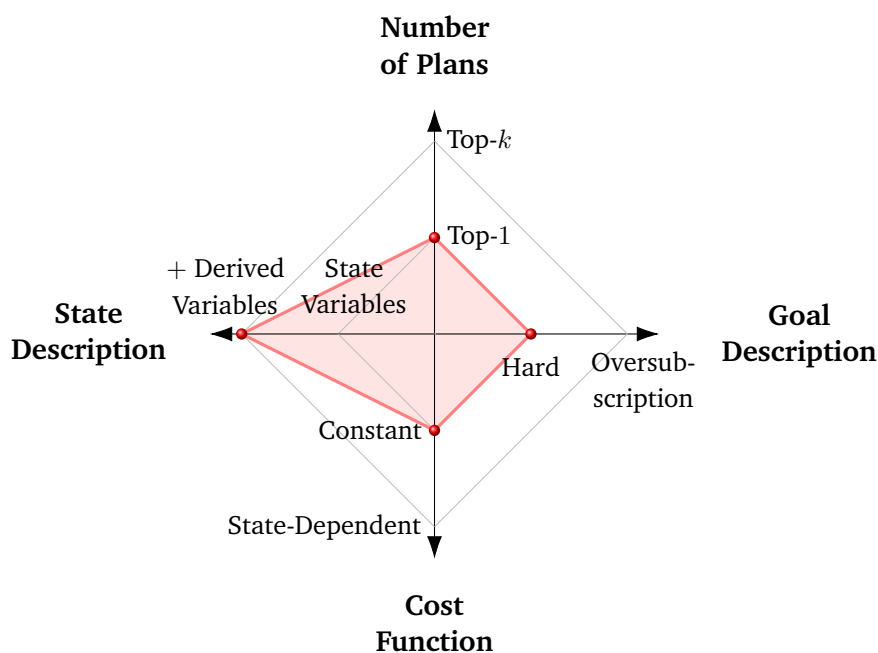


Figure 4.1: Overview of extensions for classical planning, where the red color denotes the planning formalism supported by the proposed symbolic search approach.

to cost-optimal planning, are not defined for their use with axioms. The few heuristics that support axioms are based on naive relaxations that consider axioms as zero-cost actions, which may greatly reduce the informativeness of the heuristics. One exception is the axiom-aware delete relaxation heuristic, obtained by applying a model for state constraints to planning with axioms (Haslum et al. 2018; Ivankovic and Haslum 2015). While these heuristics are often informative, they are also time-consuming to compute and therefore often do not pay off in terms of coverage or runtime.

In this chapter, we define and motivate planning with axioms and summarize the known complexity and compilability results for this setting. Then we describe three different sound and complete ways to extend symbolic search algorithms to support axioms natively introduced by Speck, Geißer, Mattmüller, and Torralba (2019) (Figure 4.1). The empirical study on different planning domains shows that the symbolic axiom encodings of Speck, Geißer, Mattmüller, and Torralba (2019) yield an optimal planner that supports axioms and compares favorably with other state-of-the-art methods.

4.1 Formalism

Planning with axioms extends the formalism of classical planning (Definitions 1 and 2) with a set of derived variables and axioms as follows (Helmert 2008; Thiébaux, Hoffmann, and Nebel 2005).

Definition 8 (Planning with Axioms). A *planning task with axioms* is a tuple $\Pi = \langle \Xi, \mathcal{I}, \mathcal{G}, \mathcal{B} \rangle$ with $\Xi = \langle \mathcal{V}, \mathcal{O}, \mathcal{C}, \mathcal{D}, \mathcal{A} \rangle$ that extends an ordinary planning task and domain with a set of *axioms* \mathcal{A} that is used to evaluate a set of *secondary* or *derived* propositional variables \mathcal{D} . Partial variable assignments such as the preconditions pre_o of operators $o \in \mathcal{O}$ and the goal condition are defined over primary and secondary variables $\mathcal{V} \cup \mathcal{D}$. A variable assignment s is called a *state* if it is defined for all variables in \mathcal{V} , and an *extended state* if it is defined for all variables in $\mathcal{V} \cup \mathcal{D}$. With \mathcal{S} we refer to the set of all states and with \mathcal{S}_E to the set of all extended states. \mathcal{A} is a finite set of axioms over $\mathcal{V} \cup \mathcal{D}$. An *axiom* has the form $h \leftarrow b$ where the head h is a variable in \mathcal{D} and the *body* b is a finite conjunction of positive or negative literals over $\mathcal{V} \cup \mathcal{D}$.

The set of axioms is partitioning into a totally ordered set of axioms layers $\mathcal{A}_1 \prec \dots \prec \mathcal{A}_i \prec \dots \prec \mathcal{A}_k$ with $i = 1, \dots, k$. The layer of an axiom is defined by the layer of its head, which is determined by a partition of the set of derived variables into subsets $\mathcal{D}_1 \prec \dots \prec \mathcal{D}_k$. We assume that this partition forms a *stratification*, i.e., that for all $i = 1, \dots, k$, and for each $d_i \in \mathcal{D}_i$, it holds that (1) if $d_j \in \mathcal{D}_j$ appears in the body of an axiom with head d_i , then $j \leq i$ and (2) if $d_j \in \mathcal{D}_j$ appears negated in the body of an axiom with head d_i , then $j < i$.

The semantics of axioms is as follows: (1) to evaluate a derived variable, only axioms in the current or previous layers have to be considered, and (2) axioms have negation-as-fault semantics, i.e., if a fact cannot be derived as true, it is assumed to be false in subsequent layers (Speck, Geißer, Mattmüller, and Torralba 2019). Given a state $s \in \mathcal{S}$ (over \mathcal{V}), the extended state $\mathcal{A}(s) \in \mathcal{S}_E$ (over $\mathcal{V} \cup \mathcal{D}$) is uniquely defined by the standard stratified semantics (Apt, Blair, and Walker 1988; Thiébaux, Hoffmann, and Nebel 2005). In other words, axioms are evaluated in a layer-by-layer fashion using fixed point computations. More precisely, given a state $s \in \mathcal{S}$, first all derived variables $d \in \mathcal{D}$ are set to their default value (false), i.e., $\neg d$. Second, a fixed point computation is performed for each axiom layer in sequence to determine the final values of the derived variables (Helmert 2008). The evaluated derived variables together with the state s form the extended state $\mathcal{A}(s)$.

Intuitively, the axioms form a background theory that makes it possible to capture/derive some properties of states, i.e., secondary variables, from the primary state variables. Consequently, an operator can only change the values of the primary variables (eff_o is a partial variable assignment over \mathcal{V}), and based on the primary variables in the successor state, the values of the secondary variables can be derived using the axioms. Moreover, an operator is applicable in a state $s \in \mathcal{S}$ if $pre_o \subseteq \mathcal{A}(s)$ and a state $s \in \mathcal{S}$ is a goal state if

$\mathcal{G} \subseteq \mathcal{A}(s)$. For a detailed description and discussion of the semantics of axioms, see Helmert (2006, 2008) and Thiébaux, Hoffmann, and Nebel (2005).

Example 4 illustrates the semantics and usefulness of axioms in the context of classical planning.

Example 4. Consider the *navigate* operators of Example 1, which navigates the rover from one cell to another with a cost of 0. This verbose modeling of the rover navigation leads to a larger state space and longer plan than necessary. Reachability can be expressed as a recursive property with axioms and derived properties. To model a *navigate* operator with axioms that moves the rover to a reachable location *loc*, we introduce a derived variable *reachable(loc)* as a precondition for the *navigate* operator. The values for the derived variables *reachable(loc)* are determined by a single layer of axioms \mathcal{A}_1 .

- $reachable(loc) \leftarrow rover-at = loc$
- $reachable(to) \leftarrow reachable(from) \wedge free(to) \wedge adjacent(from, to)$

Intuitively, the first axiom means that the current location of the rover is reachable. The second axiom means that a free location adjacent to a reachable location is also reachable. By applying the axioms until a fixed point is reached, the current state *s* is expanded to $\mathcal{A}(s)$, which then contains the information about the actual reachable locations (*reachable(loc)*) based on the current state *s*, which contains the rover location and all relevant information about the map. This encoding provides a natural and concise modeling of the transitive closure property of reachability. Moreover, it provides not only a smaller state space, but also a shorter plan $\pi^r = \langle navigate-7-1, launch-7-1, fly-7-1-to-6-1, take-img-6-1, fly-6-1-to-10-1, take-img-10-1, fly-10-1-to-7-1, land-7-1, navigate-0-5 \rangle$, which avoids the irrelevant choice of the exact path the rover must take.⁴

In Example 4, the reachable cells for the rover are always the same, but this is not necessarily true in general. For example, it could be possible that actions by the rover make some cells impassable. A similar scenario where reachability can change from state to state is Sokoban, where the locations of the boxes affect the reachability of the cells (Ivankovic and Haslum 2015; Miura and Fukunaga 2017).

4.2 Complexity and Compilability

Helmert (2008) showed that planning with axioms (Definition 8) is PSPACE-complete, since it generalizes classical planning (hardness) and a non-deterministic Turing machine can solve the bounded plan existence problem for planning with axioms with polynomial space (completeness).

⁴If the navigation operators of the rover would have non-zero costs, such a modeling would require state-dependent action costs, which we discuss in more detail in Chapter 5.




Figure 4.2: Visualization of symbolic *forward* search with the \mathcal{O} -based encoding of axioms (Speck, Geißer, Mattmüller, and Torralba 2019).

Theorem 3 (Helmert 2008). *Bounded plan existence of planning with axioms is PSPACE-complete.* \square

Thiébaux, Hoffmann, and Nebel (2005) show that axioms are an essential feature of the planning language PDDL (Hoffmann and Edelkamp 2005; McDermott et al. 1998), proving that it is in general not possible to compile away axioms without a super-polynomial growth of the plan length or description size.⁵ Furthermore, Thiébaux, Hoffmann, and Nebel (2005) argue that axioms are necessary to model real-world problems in a compact and elegant way, as they allow to model complex action preconditions (e.g. transitive closure property) or goals. From these two observations, we can conclude that it is desirable and necessary to have search algorithms that can handle axioms directly with native support to solve some real-world problems.

4.3 Symbolic Search

Speck, Geißer, Mattmüller, and Torralba (2019) introduce three different variants to support and represent axioms in symbolic search. All three encodings are sound and complete, allowing for optimal and complete symbolic search. We briefly describe the idea of the three symbolic variants and present an empirical evaluation comparing the symbolic approaches (Speck, Geißer, Mattmüller, and Torralba 2019) with explicit heuristic search approaches for planning with axioms (Ivankovic and Haslum 2015). For a more detailed explanation of the symbolic encodings, we refer to Speck, Geißer, Mattmüller, and Torralba (2019).

⁵Note that Thiébaux, Hoffmann, and Nebel (2005) also show that the 1-step planning with axioms problem is EXPTIME-complete, which results from the fact that they consider planning tasks described in PDDL, while we consider planning tasks and, in particular, axioms in a grounded and normalized form.




Figure 4.3: Visualization of symbolic *forward* search with the \mathcal{V} -based encoding of axioms (Speck, Geißer, Mattmüller, and Torralba 2019).

\mathcal{O} -Based Encoding. The most straightforward way to integrate axioms into symbolic search is the so-called \mathcal{O} -based encoding, where each axiom is interpreted as an operator, where the body is a (conditional) precondition and the head is the (conditional) effect. Unlike the original operators, the axiom operators $\mathcal{O}_{\mathcal{A}}$ affect a derived variable rather than a primary variable. With this encoding, it is possible to expand a set of states, i.e., to evaluate the derived variables of each state in a set of states in a symbolic way with a fixed point computation that represents the operators $\mathcal{O}_{\mathcal{A}}$ as a transition relation with BDDs. Figure 4.2 illustrates this procedure. The general idea is to perform this fixed point computation after each application of the actual operators to expand all newly generated successor states.

The following two axiom encodings are based on the idea of precomputing a symbolic representation over the primary variables \mathcal{V} for each derived variable $d \in \mathcal{D}$ if d is true (Definition 9). This overcomes the problem of the \mathcal{O} -based encoding, which requires an expensive fixed point computation to be performed in each planning step. Speck, Geißer, Mattmüller, and Torralba (2019) present an efficient algorithm for computing the primary representation for each derived variable using BDDs.

Definition 9 (Primary Representation). Let $d \in \mathcal{V} \cup \mathcal{D}$ be a (primary or derived) variable and \mathcal{A} a set of axioms. The *primary representation* of d is the set of states S_d , which contains all states over \mathcal{V} where d is evaluated to true, i.e., $S_d = \{s \in \mathcal{S} \mid \mathcal{A}(s) \models d\}$.

\mathcal{V} -Based Encoding. Figure 4.3 illustrates the \mathcal{V} -based encoding, where each state of a given set of successor states S is expanded. Starting from the primary representation (Definition 9), the derived variable d in a state $s \in S$ is true if $s \in S_d$. Therefore, $\chi_{\{\mathcal{A}(s) \mid s \in S\}} = \chi_S \wedge \bigwedge_{d \in \mathcal{D}} (d \leftrightarrow \chi_{S_d})$ represent the extended set of states.

Symbolic Translation. In contrast to the \mathcal{O} -based and \mathcal{V} -based encodings, the symbolic translation encoding completely omits derived variables by performing the search using a translation of the original planning task. The idea of the translation is to replace all occurrences of derived variables in the planning task with their corresponding primary representation. In particular, derived variables in operator preconditions and the goal formula are replaced by their




Figure 4.4: Visualization of symbolic *forward*, *backward* and *bidirectional* search with the symbolic translation encoding of axioms (Speck, Geißer, Mattmüller, and Torralba 2019).

corresponding primary representation (Figure 4.4). Thus, no reasoning about derived variables during the actual search is required. Note that this symbolic translation encoding is different from compilations that convert PDDL with axioms to PDDL without axioms (Gazen and Knoblock 1997; Thiébaux, Hoffmann, and Nebel 2005), since at no point is an explicit version of the compiled task (a new PDDL representation) created. While the primary representation, and thus the symbolic translation, can lead to exponential size growth in the worst case, in practice it is shown that the concise representation of BDDs alleviates this problem (see Empirical Evaluation).

Speck, Geißer, Mattmüller, and Torralba (2019) define the \mathcal{O} -based and the \mathcal{V} -based encoding for symbolic forward search only. It is an open question how to perform backward search, since it is not clear how to efficiently reverse intermediate inferences about the values of derived variables. In contrast, the symbolic compilation encoding performs all inferences over the derived variables as a preprocessing step. Thus, it allows to perform backward search and bidirectional search, which can significantly improve the planning performance.

Empirical Evaluation. Table 4.1 shows the coverage (number of optimally solved instances) on different domains with axioms modeling verification problems (Edelkamp 2003; Ghosh, Dasgupta, and Ramesh 2015), multi-agent planning with belief sets (Kominis and Geffner 2015) or elevator control (Koehler and Schuster 2000). Explicit A^* search (Hart, Nilsson, and Raphael 1968) is evaluated with the h_{blind} heuristic and the (naive) h_{max} heuristic (Ivankovic and Haslum 2015), which are the best known heuristics that are admissible and support axioms. We see that a BDD-based search is competitive with the \mathcal{O} -based and \mathcal{V} -based approaches compared to the explicit A^* search. This is mainly due to the good performance in the Miconic Axioms domain, which consists of many instances. However, symbolic search using the symbolic translation encoding proves to be the best performing and overall dominant approach, especially with bidirectional search.

Algorithm	BDD Search						
	A* (explicit)		\mathcal{O} -based	\mathcal{V} -based	Sym. Translation		
	h_{blind}	h_{max}	fwd	fwd	fwd	bwd	bid
<i>International Planning Competition (IPC)</i>							
Blocks Axioms (35)	18	18	15	15	21	18	30
Grid Axioms (5)	1	2	1	1	1	0	3
Miconic Axioms (150)	60	60	127	150	150	150	150
Optical Telegraphs (48)	2	2	2	2	4	0	4
PSR Middle (50)	35	35	32	39	50	50	50
PSR Large (50)	14	14	13	15	24	23	25
Philosophers (48)	5	5	9	9	12	4	12
<i>Complex Preconditions at IPC</i>							
Assembly (30)	0	0	6	5	9	8	11
Airport Adl (50)	19	21	14	12	20	11	19
Trucks (30)	6	8	9	9	9	4	8
<i>Ivankovic and Haslum (2015)</i>							
Blocker (7)	7	7	4	5	5	5	5
Social Planning (2)	2	2	2	2	2	2	2
Sokoban Axioms (25)	19	20	7	7	18	20	20
<i>Ghosh, Dasgupta, and Ramesh (2015)</i>							
Acc cc2 (7)	7	7	7	7	7	7	7
Grid cc2 (13)	8	7	0	0	0	0	0
<i>Kominis and Geffner (2015)</i>							
Collab And Comm (1)	1	1	0	0	0	0	0
Muddy Children (1)	1	1	1	1	1	0	1
Muddy Child (1)	1	1	1	1	1	1	1
Sum (1)	1	1	1	1	1	0	1
Word Rooms (2)	2	2	0	0	0	0	0
w/o Miconic (406)	149	154	134	131	185	153	199
Sum (556)	209	214	251	281	335	303	349

Table 4.1: Coverage (number of optimally solved instances) for explicit and symbolic search algorithms on domains with axioms (Speck, Geißer, Mattmüller, and Torralba 2019).

State-Dependent Action Costs

The fashion of the world is to avoid cost, and you encounter it.

— William Shakespeare (1598)

Core Publications of this Chapter

- **David Speck**, David Borukhson, Robert Mattmüller, and Bernhard Nebel (2021). “On the Compilability and Expressive Power of State-Dependent Action Costs”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 358–366. **(Best Student Paper Runner-Up Award)**
- **David Speck**, Florian Geißer, and Robert Mattmüller (2018a). “Symbolic Planning with Edge-Valued Multi-Valued Decision Diagrams”. In: *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*. Ed. by Mathijs de Weerd, Sven Koenig, Gabriele Röger, and Matthijs Spaan. AAAI Press, pp. 250–258. **(Partly based on ideas from my master thesis)**

In classical planning, it is common to assume constant action costs (Bäckström and Nebel 1995; Fikes and Nilsson 1971). This often results in increased effort for the modeler, because when a planning problem inherently involves state-dependent action costs (sometimes called conditional costs), the modeler of the problem must therefore distribute these costs over multiple copies of the original action. In addition, the structure of the original cost function is hidden, which, however, could provide useful information and more compact representation possibilities for planning algorithms (Geißer 2018). If we consider,

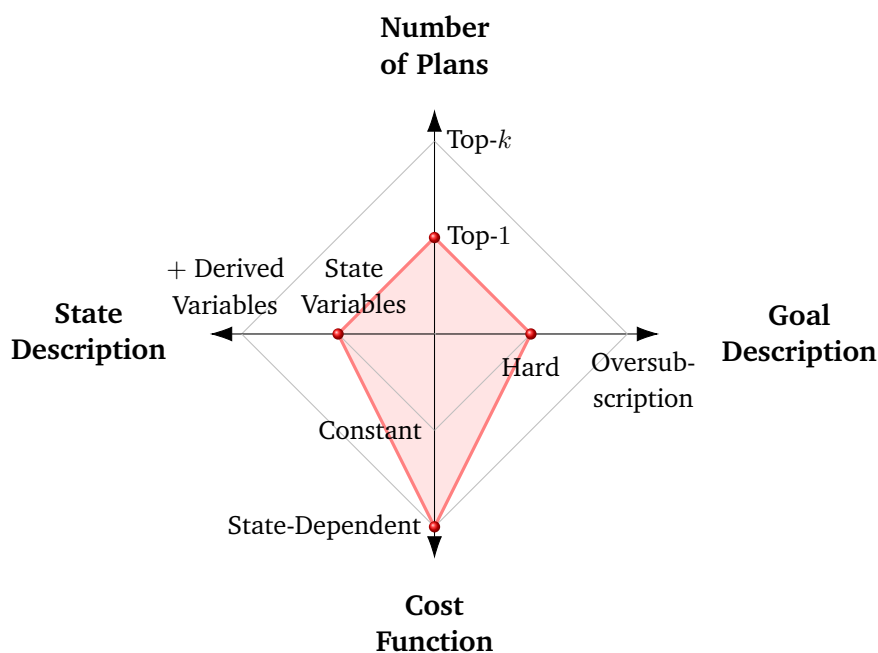


Figure 5.1: Overview of extensions for classical planning, where the red color denotes the planning formalism supported by the proposed symbolic search approach.

e.g., probabilistic planning in form of factorized Markov decision processes (Puterman 1994), state-dependent action costs/rewards have been the standard for a long time (Geißer 2018; Geißer, Speck, and Keller 2020; Sanner 2010) and are supported by many different approaches (Cui and Khardon 2018; Geißer and Speck 2018; Keller and Eyerich 2012). Therefore, recently there has been increased interest in classical planning with state-dependent action costs (Corraya et al. 2019; Drexler, Seipp, and Speck 2021; Drexler, Speck, and Mattmüller 2020; Geißer 2018; Haslum et al. 2018; Ivankovic, Gordon, and Haslum 2019; Keller and Geißer 2015; Keller, Pommerening, et al. 2016; Mattmüller et al. 2018).

In this chapter, we define and motivate planning with state-dependent action costs before summarizing the associated complexity and compilability investigated by Speck, Borukhson, et al. (2021). Then, complete and optimal symbolic search algorithms based on BDDs and EVMDDs are presented and explained for planning with state-dependent action costs (Speck, Geißer, and Mattmüller 2018a,b) (Figure 5.1). The empirical evaluation shows that the native support of state-dependent action costs within symbolic search can be beneficial overall compared to other explicit search approaches that are partially based on compilations.

5.1 Formalism

A planning domain and task with state-dependent action costs (sdac) is defined as follows (Geißer 2018; Geißer, Keller, and Mattmüller 2015).

Definition 10 (Planning with Sdac). An *sdac planning task* $\langle \Xi, \mathcal{I}, \mathcal{G}, \mathcal{B} \rangle$ with $\Xi = \langle \mathcal{V}, \mathcal{O}, (\mathcal{C}_o)_{o \in \mathcal{O}} \rangle$ is identical to a classical planning task (Definition 2), except for having (*local*) *state-dependent action cost functions* $\mathcal{C}_o : \mathcal{S} \rightarrow \mathbb{N}_0$ for each operator $o \in \mathcal{O}$. The set of all operators \mathcal{O} induces a *state-dependent action cost function* or *cost function* $\mathcal{C} : \mathcal{S} \times \mathcal{O} \rightarrow \mathbb{N}_0$ such that $\mathcal{C}(s, o) = \mathcal{C}_o(s)$ for all $s \in \mathcal{S}$. A plan $\pi = \langle o_0, \dots, o_{n-1} \rangle$ for an sdac planning task that generates a sequence of states s_0, \dots, s_n generalizes a classical plan (Definition 3) by considering for the cost computation the state in which each operator is applied, i.e., $\text{cost}(\pi) = \sum_{i=0}^{n-1} \mathcal{C}_{o_i}(s_i)$.

In general, such a state-dependent cost function can have an arbitrary form and even be uncomputable (Geißer 2018). In practice, however, it is useful to restrict the form and expressiveness of the cost function. Similar to Geißer (2018), the focus of this work is mainly on cost functions that can be evaluated in polynomial time and are in concise form (Definition 11). However, in the theoretical complexity and compilability analysis, a comprehensive classification of compilability and non-compilability for planning with sdac is provided that also considers more complex classes of cost functions (Speck, Borukhson, et al. 2021).

Definition 11 (Operator Cost Function). We define a language \mathcal{L} by the following Backus normal form:

$$t ::= c \mid v \mid t + t \mid t - t \mid t \cdot t \mid |t|,$$

where $c \in \mathbb{Z}$, $v \in \mathcal{V}$. The semantic of \mathcal{L} is defined over states $s \in \mathcal{S}$ as follows:

$$\begin{aligned} \mathcal{C}^c(s) &= c \\ \mathcal{C}^v(s) &= s(v) \\ \mathcal{C}^{t \circ t'}(s) &= \mathcal{C}^t(s) \circ \mathcal{C}^{t'}(s) \quad \text{for } \circ \in \{+, -, \cdot\} \\ \mathcal{C}^{|t|}(s) &= |\mathcal{C}^t(s)| \end{aligned}$$

For a given term $t \in \mathcal{L}$, the interpretation \mathcal{C}^t specifies the operator cost function, where we restrict the allowed operator cost function to a positive range, i.e., $\mathcal{C}^t : \mathcal{S} \rightarrow \mathbb{N}_0$. In the following, we often identify a cost function \mathcal{C}^t with the term t that defines it.

We emphasize that planning with sdac is a generalization of classical planning. Classical planning is an important special case in which we have a constant action cost, i.e., operator costs that are independent of the state

in which the operator is applied. Modeling operator costs as state-dependent allows for a more natural and concise representation of planning problems as Example 5 illustrates.

Example 5. Consider the drone flight operators of Example 1. In classical planning (Definition 2), we need one *fly* operator for each pair of cells *from* = (x, y) and *to* = (x', y') to model the Manhattan distance as constant costs. Note that such modeling is often used to represent motion operators with costs in the International Planning Competition (e.g., TRANSPORT domain). With *sdac*, we are able to model such operators in a natural and concise way by specifying the Manhattan distance directly as the cost function of the *fly* operator. For each cell *to* = (x', y') , we specify an operator *fly-to- x' - y'* that is applicable when the drone is launched and has the effect of the drone being at (x', y') . The cost of the operator *fly-to- x' - y'* is $|x_{drone} - x'_{drone}| + |y_{drone} - y'_{drone}|$.

5.2 Complexity and Compilability

The work of Speck, Borukhson, et al. (2021) addresses the computational complexity and compilability of planning with *sdac*. They show that planning with *sdac* has the same complexity as ordinary planning, i.e., is PSPACE-complete. It is natural that a cost function should not be computationally more difficult than planning itself, i.e., more difficult than PSPACE. Otherwise, the evaluation of the cost function would dominate the computation instead of solving the actual planning task, which should be the actual challenge. By restricting the complexity of cost functions to FPSPACE, we capture cost functions that can be evaluated in polynomial time (Definition 11), up to more complex cost functions such as the cost of a robot motion operation (LaValle 2006; Reif 1979).

Theorem 4 (Speck, Borukhson, et al. 2021). *Bounded plan existence of planning with *sdac* is PSPACE-complete, if the cost function is in FPSPACE.* \square

Recall that the same computational complexity of two planning formalisms does not imply the same expressive power (Nebel 2000). While existing translations (EVMDD-based and combinatorial (Geißer 2018; Geißer, Keller, and Mattmüller 2015)) of *sdac* planning tasks into classical (constant cost) planning tasks work well in practice, both can lead to an exponential increase in task size, which means that strictly speaking they are not compilations by Definition 4. Speck, Borukhson, et al. (2021), however, analyzed the compilability of *sdac* and obtained results of possibility and impossibility that depend on the desired plan length and the complexity of the cost function. For this purpose, it is assumed that the cost function is provided in a concise form of a deterministic Turing machine (DTM). It should be noted that, in general, other concise forms of cost functions, such as computer programs or the format

		Desired Plan Length Preservation		
		linearly	polynomially	don't care
c Complexity	FP	<i>impossible</i>	<i>possible</i> using DTM compilation	
	FPSPACE	<i>impossible</i> (unless polynomial hierarchy collapses at the third level)		<i>possible</i> using DTM compilation

Table 5.1: Existence results for compilation schemes preserving task sizes polynomially depending on the computational complexity of the cost functions (rows) and the desired plan length preservation (columns) (Speck, Borukhson, et al. 2021).

defined in Definition 11, can be simply converted to this form by specifying a DTM that computes/simulates the corresponding cost function. Table 5.1 gives an overview of the compilation results, which we discuss below.

Possibility Results. It can be shown that there exists a compilation scheme that preserves the plan length polynomially when the cost function is in FP. Moreover, there is a valid compilation scheme (unbounded plan length) if the cost function is in FPSPACE. The underlying idea of the compilation scheme that yields these two results is to simulate a DTM that computes the cost function $\mathcal{C}_o(s)$ of the operator $o \in \mathcal{O}$ for a current state $s \in \mathcal{S}$ within the planning task. For this purpose, each operator is compiled into an operator sequence that ensures that the task size and plan length are bounded by the space and time complexity of the corresponding DTM used to compute the cost function (Speck, Borukhson, et al. 2021). Note that this compilation makes it possible to encode various different cost functions that can be computed with a DTM, although the usefulness of the resulting planning task obviously depends heavily on the complexity of the DTM.

Impossibility Results. Finally, Speck, Borukhson, et al. (2021) proved that it is impossible to compile away, i.e., to find a compilation scheme that preserves plan length linearly when the cost function is in FP. Moreover, it is impossible to preserve the plan length polynomially when the cost function is in FPSPACE, unless the polynomial hierarchy collapses to the third level.

Considering these results on complexity and compilability, one can see that it may be helpful not to compile sdac at all, but to keep it in the model and support it natively in algorithms, which avoids the overhead introduced by compilation (Speck, Borukhson, et al. 2021). In particular, the fact that cost functions in FP, such as the fragment considered in practice (Definition 11),

require polynomial blowup in terms of domain description size and original plan length may make the use of the compilation technique infeasible. In the following, we show how it is possible to natively support and represent sdac with decision diagrams to perform a symbolic search.

5.3 Symbolic Search

Speck, Geißer, and Mattmüller (2018a) introduce a complete and optimal approach that performs symbolic search with EVMDDs. In contrast to symbolic search with BDDs, where costs are represented by partitioning sets of states into subsets with identical cost values (see Example 2), with EVMDDs it is possible to represent a set of reachable states together with reachability costs simultaneously. The latter is possible by assigning to reachable states $S \subseteq \mathcal{S}$ the corresponding cost $g \in \mathbb{N}_0$ and to unreachable states the cost ∞ . Speck, Geißer, and Mattmüller (2018a) introduced all the necessary operations such as the image and preimage operations for EVMDDs to enable symbolic search. The main advantages of using EVBDDs/EVMDDs over ADDs and BDDs are that they can represent certain functions exponentially more compactly and that encoding reachability and cost in the open list, closed list, and transition relations simultaneously can also result in a more concise representation (Speck, Geißer, and Mattmüller 2018a).

In addition to the EVMDD-based symbolic search for sdac planning tasks, we propose and study in this thesis a novel variant based on BDDs. The underlying idea is to first create an ADD (instead of an EVMDD) that represents each operator as a transition relation that includes costs, before partitioning the ADD into multiple BDDs. In other words, we create multiple transition relations for each operator with sdac, one for each possible cost value, encoding as a precondition the corresponding condition for the cost. This approach may result in multiple transition relations with different costs, but allows us to perform the classical symbolic search with cost bucketing, using the full power of sophisticated BDD operations and libraries.

Note that the representation of the cost function as decision diagram can lead to an exponential increase in the worst case. However, as usual with decision diagrams, it is assumed that the representation size of the cost functions is manageable.

Example 6. Consider an operator $o = \langle \neg x, x \rangle$ with sdac $\mathcal{C}_o(s) = 5y + 1$ for $s \in \mathcal{S}$. Figure 5.2 visualizes the transition relation of o represented with different decision diagrams, where the precondition (predecessor states) are encoded with unprimed variables and the effects (successor states) with primed variables. Figure 5.2a represents the EVMDD representing the transition relation of o , while encoding the cost $\mathcal{C}_o(s)$ for each state $s \in \mathcal{S}$ at the same time. To represent the operator o with BDDs, multiple BDDs are necessary to decompose the cost function. Figure 5.2b illustrates two BDDs representing the

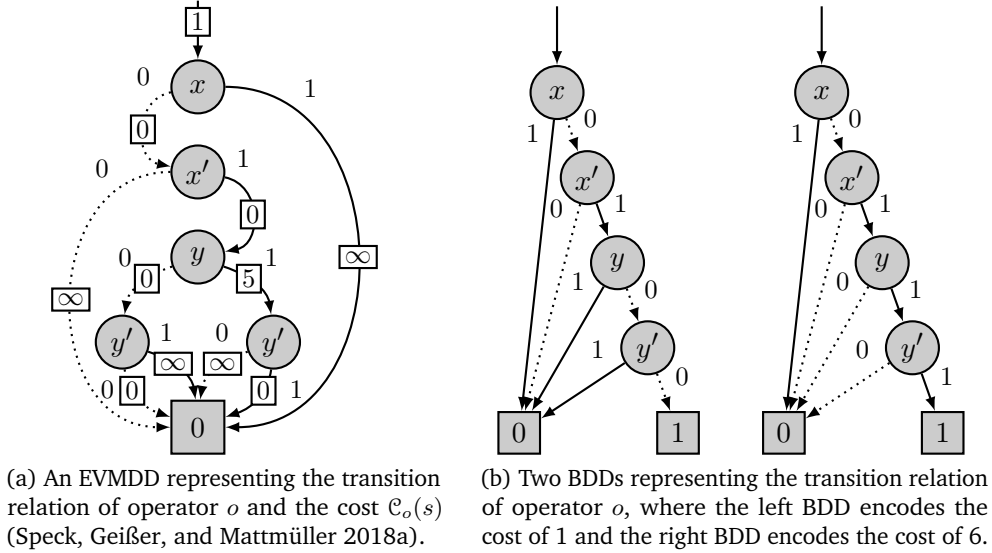


Figure 5.2: Visualization of different variants for a symbolic representation of an operator $o = \langle \neg x, x \rangle$ with state-dependent action costs $\mathcal{C}_o(s) = 5y + 1$.

two possible cost values of $\mathcal{C}_o(s)$, specifically 1 and 6. The left BDD encodes the cost of 1 by encoding $\neg y$ in the precondition. Note that by adding the new precondition $\neg y$, $\neg y$ also holds in all successor states, i.e., $\neg y'$, since variable y does not occur in the effect (closed world assumption). Similarly, the right BDD encodes y as a precondition to obtain a cost of 6.

Empirical Evaluation. While Speck, Geißer, and Mattmüller (2018a) reported results comparing symbolic search with EVMDDs and translation-based approaches, here we present new empirical results incorporating recent advances, planners, and benchmarks for planning with sdac (Corraya et al. 2019; Geißer 2018; Speck, Borukhson, et al. 2021).

Table 5.2 shows the coverage (number of optimally solved instances) of explicit A* search (Hart, Nilsson, and Raphael 1968) with various translation-based heuristics (Geißer 2018) and forward, backward and bidirectional symbolic search with all data structures represented as either EVMDDs (Speck, Geißer, and Mattmüller 2018a) or BDDs. In explicit A* search for planning with sdac, the cost function is represented as an EVMDD, which is used to evaluate the actual costs (g -values) (Geißer 2018). We consider the h_{blind} heuristic and the h_{max} heuristic (Bonet and Geffner 1999), which is analyzed in two versions: h_{max}^c uses the combinatorial translation (Geißer 2018) and $h_{\text{max}}^{\text{dd}}$ uses the EVMDD translation (Geißer, Keller, and Mattmüller 2015) to compute the heuristic values.

Algorithm Domain (# Tasks)	A* (explicit)			EVMDD Search			BDD Search		
	h_{blind}	h_{max}^c	$h_{\text{max}}^{\text{dd}}$	fwd	bwd	bid	fwd	bwd	bid
Asterix (30)	10	5	10	30	29	30	30	28	30
Pegsol (100)	96	0	96	86	16	88	96	23	96
Gripper (30)	8	4	8	11	8	12	20	14	19
PSR Large (100)	33	12	39	36	20	36	28	23	28
PSR Middle (100)	79	34	84	88	55	90	74	58	74
Openstacks (70)	9	10	10	40	30	38	54	46	51
Transport (30)	24	19	24	24	23	29	24	23	28
TSP (26)	21	12	21	15	10	13	18	10	18
Sum (486)	280	96	292	330	191	336	344	225	344

Table 5.2: Coverage (number of optimally solved instances) for explicit and symbolic search algorithms on versions of planning domains with state-dependent action costs.

Overall, it can be seen that native support for sdac within symbolic search compares favorably with explicit search approaches. Comparing symbolic search with BDDs and EVMDDs, there are domain-wise differences in coverage, with BDD-based symbolic search performing best overall. However, the structural advantages of EVMDDs pay off especially in the PSR domain, where the cost function is complex and contains derived variables represented with the symbolic translation described in Chapter 4.⁶ Looking at the explicit approaches, we can see that in many domains it is not possible to perform a combinatorial translation h_{max}^c , which often requires exponentially many operators and thus solves the least number of tasks. Explicit A* search with the h_{blind} and the $h_{\text{max}}^{\text{dd}}$ heuristics performs much better. These configurations use EVMDDs to concisely represent the cost functions and as a basis for the heuristic computation, which seems to pay off. It turns out that the heuristic estimates of $h_{\text{max}}^{\text{dd}}$ can help the search, but the performance of $h_{\text{max}}^{\text{dd}}$ is still often not comparable to symbolic blind search. However, there are domains where explicit heuristic search performs better than symbolic search. This shows that, similar to classical planning, a potential portfolio planner of these complementary search strategies could lead to a state-of-the-art planner that combines both strengths (Sievers et al. 2019).

⁶In Section 8.1, the relationship and interplay of the different planning extensions is discussed in more detail.

Oversubscription Planning

A goal is not always meant to be reached, it often serves simply as something to aim at.

— Bruce Lee

Core Publication of this Chapter

- **David Speck** and Michael Katz (2021). “Symbolic Search for Oversubscription Planning”. In: *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*. Ed. by Kevin Leyton-Brown and Mausam. AAAI Press, pp. 11972–11980.

In conventional classical planning, goal states are defined by a goal condition, which is treated as a hard constraint that must be fully satisfied by any solution/plan. In many real-world applications, however, the description of goals is oversubscribed, i.e., there are a large number of desirable, often competing goals of varying value, and a system (e.g., a Mars rover) may not be able to achieve all of them with the available resources (e.g., battery power). In such scenarios, it is natural to assign different utility values to states and to search for feasible and achievable states that maximize the overall utility value. This planning formalism is called *partial satisfaction planning*. Here, we distinguish between *net-benefit planning* (van den Briel et al. 2004), where operator costs and state utility values, i.e., solution costs and solution utilities, are comparable, and *oversubscription planning* (Smith 2004), where those are not comparable.

Symbolic search in the context of partial satisfaction planning has been studied to a limited extent. For net-benefit planning, symbolic branch-and-bound search was introduced (Edelkamp and Kissmann 2009; Kissmann 2012), which considers so-called soft-goal planning tasks, i.e., planning tasks in which

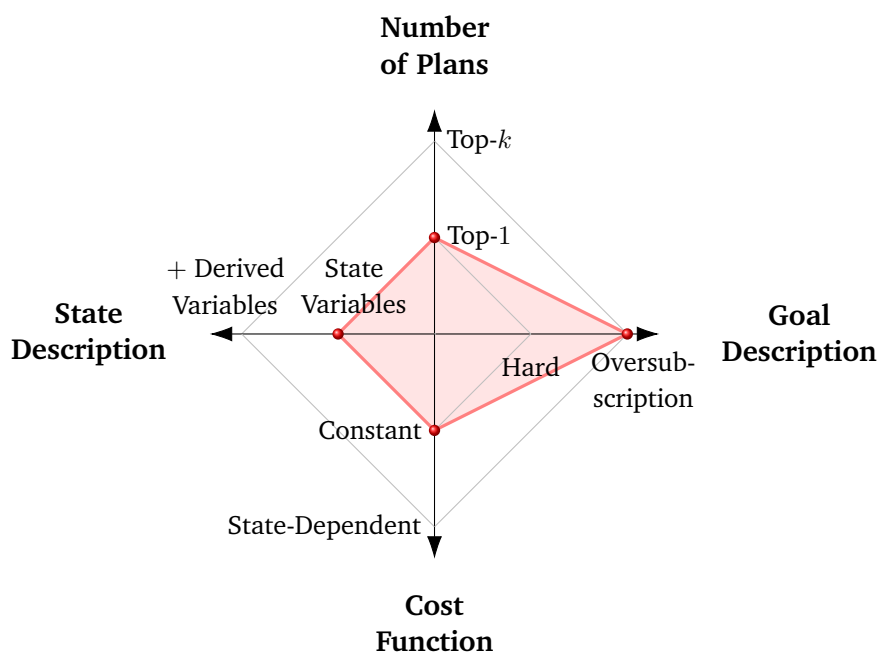


Figure 6.1: Overview of extensions for classical planning, where the red color denotes the planning formalism supported by the proposed symbolic search approach.

the utility function is a weighted sum over a subset of the state variables. Since Keyder and Geffner (2009) provided a simple and concise compilation from net-benefit planning (soft-goal planning tasks) to classical planning, which is most commonly used in practice, we will focus on oversubscription planning, which is considered more challenging (Domshlak and Mirkis 2015). For oversubscription planning, there are a variety of explicit (heuristic) search approaches, mainly based on adapting existing heuristics from classical planning to oversubscription planning (Domshlak and Mirkis 2015; García Olaya, Rosa, and Borrajo 2021; Katz, Keyder, Pommerening, et al. 2019a; Mirkis and Domshlak 2013; Muller and Karpas 2018). An important example of symbolic search in the context of oversubscription planning is the work of Eifler et al. (2020), in which plan explanations were investigated considering plan utilities. The work of Speck and Katz (2021) is the first to explicitly define and study in detail symbolic search for solving oversubscription planning problems.

In this chapter, we define and motivate oversubscription planning that allows to extend the goal specification (Figure 6.1), and discuss and summarize complexity and compilability results for this setting. Then, an optimal and complete symbolic search approach for oversubscription planning introduced by Speck and Katz (2021) is presented. The empirical analysis shows that

the presented symbolic approach competes favorably with explicit heuristic state-space search.

6.1 Formalism

We consider oversubscription planning tasks (Katz, Keyder, Pommerening, et al. 2019a; Smith 2004), where the solution cost and utility are *not comparable*.

Definition 12 (Oversubscription Planning). An *oversubscription planning (osp) task* is a tuple $\Pi = \langle \Xi, \mathcal{I}, \mathcal{G}, \mathcal{B} \rangle$ with $\Xi = \langle \mathcal{V}, \mathcal{O}, \mathcal{C}, \mathcal{U} \rangle$ that extends an ordinary planning task and domain by a state utility function $\mathcal{U} : \mathcal{S} \rightarrow \mathbb{N}_0$ that assigns utility to states.

Similar to planning with sdac, where we focused mainly on cost functions that can be evaluated in polynomial time, here we focus mainly on utility functions that can be evaluated in polynomial time. We assume that the utility function \mathcal{U} has the same form as an operator cost function (Definition 11).

Since we consider utility in osp, it is natural to search for a plan that maximizes utility, i.e., leads to a state with high utility within the cost bound \mathcal{B} . For this purpose, we define an osp plan as follows and introduce the concept of utility-optimal plans.

Definition 13 (Osp Plan). An *(osp) plan* π is defined as an ordinary classical plan (Definition 3), i.e., an applicable sequence of operators that leads from an initial state \mathcal{I} to a goal state $s_* \in S_*$, where the plan cost (cumulative operator cost) is within the cost limit $cost(\pi) \leq \mathcal{B}$.

The *utility* $\mathcal{U}(\pi)$ of the plan π is defined by the utility of the final state s_* induced by π , i.e., $\mathcal{U}(\pi) = \mathcal{U}(s_*)$.

Such a plan π is called *utility-optimal*, or simply *optimal* (if it is clear from the context), if there is no other plan π' for which $\mathcal{U}(\pi') > \mathcal{U}(\pi)$. A plan π is *cheapest utility-optimal* if it is a cheapest among utility-optimal plans, i.e., there is no utility-optimal plan π' for which $cost(\pi') < cost(\pi)$.

The underlying idea of osp is to model many different possible competing goals with different associated utilities. The overall goal is to achieve as much as possible given a cost \mathcal{B} that models a limited resource together with operator costs. Therefore, in practice, the “hard” goals \mathcal{G} are often underspecified or even nonexistent, as the goals can be modeled with utility (Katz, Keyder, Pommerening, et al. 2019a,b). Since all the approaches and results presented in this chapter can straightforwardly support hard goals \mathcal{G} , we consider them for completeness. As a result of this modeling, there are often many valid plans that vary strongly in their utility. Thus, the main goal and challenge of oversubscription planning is to identify a plan with high utility. The search for an utility-optimal plan, i.e., a plan with the highest utility, is called *optimal oversubscription planning* and is the main topic of this chapter. Originally, osp

was introduced to address NASA missions such as planning science experiments for a Mars rover (Smith 2004), where there are often limited resources and several different and competing goals.

Example 7. Recall Example 1, where we specified the goal as taking images at (6,1) and (10,1) and traveling equipped with the drone to “Three Forks” at (0,5). It is reasonable to assume that the trip to (0,5) with the drone is the (hard) goal \mathcal{G} , since the journey of the rover is expected to continue from there. Let us also assume that the drone has a limited battery capacity, as in reality. With the operator cost \mathcal{C} we describe the battery consumption of the drone for the respective actions, and the final plan is constrained by the battery capacity of say 20, which is represented as the cost limit $\mathcal{B} = 20$. Recall that in our model, only the operators associated with the drone have non-zero costs.

With osp, we can now express that we are more interested in the location (10,1) by assigning a utility of 10 to the states in which an image was taken at (6,1) and a utility of 15 to the states in which an image was taken at (10,1). If we consider the original plan π^r with a cost of $cost(\pi^r) = 24$ by taking images at (6,1) and (10,1), with a utility of $10 + 15 = 25$, we see that it is not feasible because it is too expensive at a cost of $cost(\pi^r) = 24 \not\leq 20 = \mathcal{B}$.

Due to the cost bound, in a utility-optimal osp plan, only the objective of taking an image at (10,1) can be satisfied, ignoring location (6,1). The final utility-optimal osp plan is as follows $\pi^r = \langle navigate-7-2, navigate-7-1, launch-7-1, fly-7-1-to-10-1, take-img-10-1, fly-10-1-to-7-1, land-7-1, navigate-7-2, \dots, navigate-0-5 \rangle$ with a cost of $cost(\pi^r) = 18$ and a utility of $\mathcal{U}(\pi^r) = 15$. If the cost limit were even lower, the plan would change to taking an image at (6,1) or even not taking any images at all.

6.2 Complexity and Compilability

In net-benefit planning, extensive complexity analyses have been proposed, showing, among other things, that it is PSPACE-complete (Aghighi and Bäckström 2015; Aghighi and Jonsson 2014). Moreover, for net-benefit planning, where it is common to define the state utility function as a weighted sum over a subset of the state variables, it was shown that these have no expressive power and can be easily compiled away (Keyder and Geffner 2009).

However, in osp, where solution cost and solution utility are *not comparable*, there is little work on complexity and compilability (Katz and Mirkis 2016). Katz and Mirkis (2016) state that optimal oversubscription planning is PSPACE-complete. There is no explicit proof of this statement in the literature, so we provide such a proof in this thesis. For this purpose, we define the bounded utility plan existence problem as follows.

Definition 14 (Bounded Utility Plan Existence). The *bounded utility plan existence* problem is the decision problem of determining whether there exists a plan with utility greater than or equal to $u \in \mathbb{N}_0$ for a given osp task Π .

Similar to planning with state-dependent action costs (Definition 10), we assume a utility function that is computationally no more difficult than planning itself, i.e., lies in FPSPACE.

Theorem 5. *Bounded utility plan existence is PSPACE-complete, if the utility function is in FPSPACE.*

Proof. PSPACE-hardness results from reducing the well-known bounded plan existence (with unit cost) problem (Theorem 4; Bylander 1994) to our problem, where the plan length is equal to the plan cost. PSPACE membership can be proved by defining a nondeterministic Turing machine (NTM) that starts with the initial state and guesses an operator to apply at each step. In addition, the NTM evaluates the utility of the current state s in each step. The NTM ends with “Yes” if $s \in S_*$ and $\mathcal{U}(s) \geq u$ and with “No” if the selected operator is not applicable or the cost bound is exceeded. Since at any point in time only the current state, the summed costs, and the computation of the utility function (in FPSPACE) need to be maintained, this NTM is in NPSPACE, which is known to be equivalent to PSPACE (Savitch 1970). \square

Note that this proof holds even if the osp formalism does not contain a goal specification \mathcal{G} , as is often the case in the literature. The PSPACE-hardness remains because we can simulate \mathcal{G} by assigning non-zero utility only to the goal states S_* , while requiring a plan with utility greater than 0, i.e., $u > 0$.

Similar to all other extensions, we can observe that the computational complexity is the same for classical and oversubscription planning. The compilability of osp to classical planning is an open research question. Whether and to what extent osp is more expressive than classical planning is not known yet. However, there are some results and arguments in the literature that suggest that osp is (significantly) more expressive than classical planning.

Katz and Mirkis (2016) investigated the computational complexity of osp given common classes of utility functions. They showed that the complexity of several of these fragments is computationally more difficult for osp than for net-benefit planning, where operator costs and state utilities are comparable, indicating that osp is computationally more challenging. Domshlak and Mirkis (2015) argue that unlike classical planning or net-benefit planning, osp does not seem to be reducible to a shortest path problem with a single source and destination. For this reason, solving osp tasks requires (1) exhaustive search algorithms and (2) sophisticated heuristics used in classical planning cannot be applied directly. Based on these observations, Speck and Katz (2021) proposed to use symbolic search for osp because symbolic search is known to be an efficient search strategy that (1) can exhaustively explore the state space and (2) does not rely on heuristics, i.e., performs a blind search.

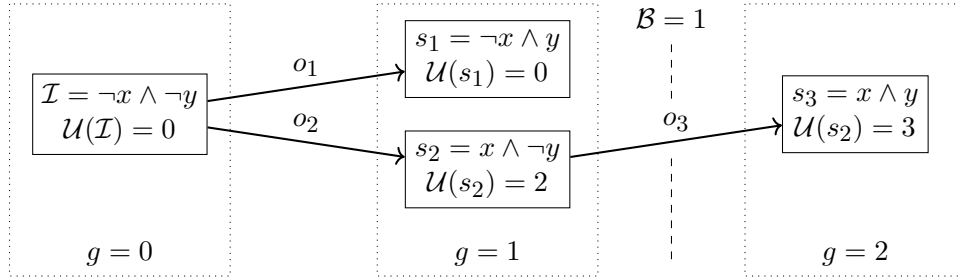


Figure 6.2: Visualization of the transition system induced by the oversubscription planning tasks Π in Example 8 (Speck and Katz 2021).

6.3 Symbolic Search

Speck and Katz (2021) define symbolic *forward* search for osp in a straightforward way and showed that their approach is complete and optimal. In fact, the presented symbolic search approach finds not only a utility-optimal plan but also a cheapest utility-optimal plan.

The underlying idea is to perform an exhaustive symbolic forward search, just as in classical planning, until the cost bound \mathcal{B} is exceeded. The utility function \mathcal{U} is a numerical function that can be represented with one ADD or several BDDs, as explained in Example 2. In each expansion step, the utility of the expanded states is symbolically evaluated and the states with the highest utility so far are maintained. In addition, representing the utility as ADDs (or BDDs) has the advantage that the highest possible value of the utility function is known directly, since it is simply the terminal node with the highest value. Finally, we stop when either all reachable states have been expanded within the cost bound or a state with the overall maximum utility has been found, since there can be no better plan. Example 8 illustrates how the symbolic approach of Speck and Katz (2021) works.

Example 8 (Speck and Katz 2021). Consider a unit-cost oversubscription planning task $\Pi = \langle \mathcal{V}, \mathcal{O}, \mathcal{C}, \mathcal{U}, \mathcal{I}, \mathcal{G}, \mathcal{B} \rangle$ with an empty goal $\mathcal{G} = \{\}$ such that all states are goal states, i.e., $\mathcal{S} = \mathcal{S}_*$. Moreover, Π has two binary variables $\mathcal{V} = \{x, y\}$, where $D_x = D_y = \{0, 1\}$, an initial state $\mathcal{I}(x) = \mathcal{I}(y) = 0$, a utility function $\mathcal{U} = 2x + xy$, and a cost bound $\mathcal{B} = 1$. There exist three operators $\mathcal{O} = \{o_1, o_2, o_3\}$, where $o_1 = \langle \neg x \wedge \neg y, y \rangle$, $o_2 = \langle \neg x \wedge \neg y, x \rangle$ and $o_3 = \langle x \wedge \neg y, y \rangle$. The induced transition system of Π is depicted in Figure 6.2.

Symbolic forward search for osp (Speck and Katz 2021) starts with a BDD representing the set of states $open = \{\mathcal{I}\} = \neg x \wedge \neg y$ consisting of a single state, namely the initial state. Figure 2.2a visualizes the ADD and Figures 2.2c to 2.2e visualize the BDDs representing the utility function \mathcal{U} with a maximum utility of 3. In the first step, we expand $\{\mathcal{I}\}$, which also forms the set of best

Algorithm	BDD Search		A_{uADD}^*		A_{mc}^*		BnB	
	uBDD	uADD	h_{blind}	h_{blind}	h_{max}^b	$h_{\text{m\&s}}^b$	h_{blind}	$h_{\text{mcut}}^{\text{mc}}$
25% Bound (1667)	1271	1274	1165	1197	1190	1074	1183	1151
50% Bound (1667)	990	993	860	901	902	828	893	867
75% Bound (1667)	866	862	718	758	738	734	735	702
100% Bound (1667)	802	793	629	668	655	676	643	618
Sum (6668)	3929	3922	3372	3524	3485	3312	3454	3338

Table 6.1: Coverage (number of optimally solved instances) for explicit and symbolic search algorithms on oversubscription planning domains (Speck and Katz 2021).

states seen so far with utility $\mathcal{U}(\mathcal{I}) = 0$. The expansion leads to two new states $open = \{s_1, s_2\} = (\neg x \wedge y) \vee (x \wedge \neg y)$, both of which can be achieved with cost $g = 1$. Since s_2 has utility $\mathcal{U}(s_2) = 2$, the set of best states changes to $\{s_2\}$. Finally, the plan $\langle o_2 \rangle$ leading to one of the best states (here s_2) is reconstructed as the cost limit of $\mathcal{B} = 1$ is exceeded. Recall that all these computations are performed with decision diagrams.

Finally, Speck and Katz (2021) argue that it is not straightforward to efficiently apply regression to osp, which would enable symbolic backward search and, in particular, symbolic bidirectional search. Unlike classical planning, where the starting point of symbolic backward search is obvious, namely the goal states represented as a compact goal formula, the starting point of symbolic backward search for osp is not obvious. The reason is that the objective is to find a final state with maximum utility, and it is not clear how this can be represented in backward search. It might be possible to partition all goal states by their utility values and represent them as separate BDDs. However, since the goal formula in osp tasks is often unspecified or even empty, this may require a regression of the entire state space with multiple backward searches. All in all, it is not clear how to efficiently apply symbolic regression to osp, which is an open research question.

Empirical Evaluation. Table 6.1 shows the coverage on four different osp benchmark sets (Speck and Katz 2021). Each benchmark set consists of 57 domains and originates from the optimal track of the IPC 1998-2014, where the goal facts are no longer hard goal facts, but have assigned utilities (Katz, Keyder, Pommerening, et al. 2019a,b). The sets differ in the cost bounds, which for each planning instance are set at 25%, 50%, 75%, or 100% of the cost of the optimal or best known solution. The two symbolic search approaches (BDD search) perform a forward search using BDDs, but differ in how the utility of a set of states is evaluated, either with multiple BDDs (uBDD) or




Figure 6.3: Coverage (number of optimally solved instances) over time for explicit and symbolic search algorithms on all oversubscription planning domains (25%, 50%, 75%, or 100% Bound) (Speck and Katz 2021).

one ADD (uADD). As a result, symbolic search with a decomposed utility function in multiple BDDs (BDD search + uBDD) works best overall. The main advantage of representing utility values with BDDs compared to ADDs is that the underlying decision diagram library CUDD (Somenzi 2015) uses techniques such as complement edges to store BDDs more compactly (Brace, Rudell, and Bryant 1990).

Overall, the described symbolic search approach (Speck and Katz 2021) performs better than explicit search approaches. Comparing symbolic search with its explicit counterpart, A_{uADD}^* search with the blind heuristic h_{blind} , we see a difference in performance due to the concise representation of state sets. Considering other publicly available planning algorithms for osp, we find that branch-and-bound search BnB (Katz, Keyder, Pommerening, et al. 2019a) and A_{mc}^* search with multiple cost functions (mc) (Katz and Keyder 2019) perform worse than symbolic search in terms of overall coverage. The performance difference between symbolic search and explicit heuristic search increases as the bound increases and thus the plan length increases. The natural explanation for this is that heuristics do not pay off, since in osp it seems rather difficult to define a heuristic that is both informative and fast to compute. This observation is underlined by Figure 6.3, which depicts the number of instances solved over time, where symbolic search dominates all other approaches after a short time. Finally, as usual, there are domains where explicit heuristic search performs better than symbolic search.

Top-k Planning

If plan A fails, remember there are 25 more letters.

— Chris Guillebeau

Core Publication of this Chapter

- **David Speck**, Robert Mattmüller, and Bernhard Nebel (2020). “Symbolic Top-k Planning”. In: *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*. Ed. by Vincent Conitzer and Fei Sha. AAAI Press, pp. 9967–9974.

In planning, it is common to assume that the model is fully specified and that the objective is to find a single plan. While in some cases a single plan may be sufficient, in practice it is often better to have many good alternative plans (Nguyen et al. 2012). The possibility of obtaining multiple high-quality plans can be of great importance in various fields and applications, such as task and motion planning (LaValle 2006; Lozano-Pérez and Kaelbling 2014; Ren, Chalvatzaki, and Peters 2021), diverse planning (Katz, Sohrabi, and Udrea 2020), scenario planning (Sohrabi, Riabov, Katz, et al. 2018), goal recognition (Sohrabi, Riabov, and Udrea 2016) or planning with user preferences (Nguyen et al. 2012; Seimetz, Eifler, and Hoffmann 2021).

The problem of determining k shortest paths for a given graph is a well-studied topic, going back to 1957 (Bock, Kanter, and Haynes 1957). However, the problem of determining k cheapest plans, known as top- k planning, has only been studied more recently (Riabov, Sohrabi, and Udrea 2014). In addition to generating a set of possible solutions for all the above mentioned applications, the systematic enumeration of plans in an anytime behavior enables the realization of a “generate-and-test framework” by searching for

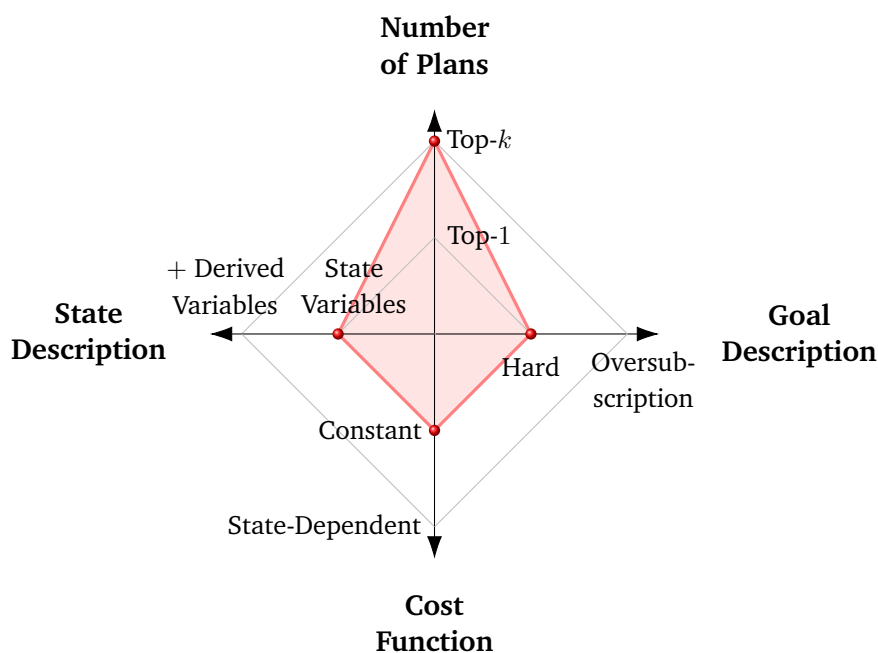


Figure 7.1: Overview of extensions for classical planning, where the red color denotes the planning formalism supported by the proposed symbolic search approach.

plans that must satisfy certain complex requirements. Thus, top- k planning can also be used to search in a simplified version of a given problem, e.g., in an overapproximated, abstracted, or decomposed model representation (Höller 2021).

In this chapter, we define and motivate top- k planning (Figure 7.1) before presenting the work of Speck, Mattmüller, and Nebel (2020) on this topic. On the theoretical side, we discuss the question of whether top- k planning is computationally more difficult than ordinary planning. On the practical side, we describe a sound and complete symbolic search approach to top- k planning. The empirical evaluation shows that symbolic search performs better than other state-of-the-art search methods for both a small and a large number of desired plans k .

7.1 Formalism

The objective of top- k planning is to determine a set of k different plans with lowest cost for a given classical planning task (Definition 2). We formally define top- k planning as follows (Katz, Sohrabi, Udrea, and Winterer 2018; Riabov, Sohrabi, and Udrea 2014; Speck, Mattmüller, and Nebel 2020).

Definition 15 (Top- k Planning). Given a planning task Π and a natural number $k \in \mathbb{N} \cup \{\infty\}$, top- k planning is the problem of determining a set of plans $P \subseteq P_{\Pi}$ such that:

1. there exists no plan $\pi' \in P_{\Pi}$ with $\pi' \notin P$ that is cheaper than some plan $\pi \in P$, and
2. $|P| = k$ if $|P_{\Pi}| \geq k$, and $|P| = |P_{\Pi}|$, otherwise.

We emphasize that top- k planning is a generalization of classical planning, where $k = 1$ corresponds to cost-optimal classical planning.

There can be infinitely many plans, since according to Definition 15 plans with cycles, i.e., plans which visit the same state multiple times, are allowed. Moreover, there can be infinitely many plans with optimal costs if cycles with 0 costs exist. In general, top- k planners not only provide a way to generate the k best plans, but also to enumerate all possible plans if no 0-actions or, more precisely, no 0-cost cycles exist. Depending on the application, plans with cycles can be important to provide all possibilities, e.g., when user preferences are not fully known or the model specification is not complete. In addition, the sequential generation of all plans allows the realization of complete search strategies that search in a simplified version of a problem at hand (Höller 2021). The symbolic approach presented in this chapter can be modified to generate only simple plans, i.e., plans without cycles (von Tschammer, Mattmüller, and Speck 2022). However, we will focus on “ordinary” top- k planning, which includes all possible plans. Example 9 illustrates the idea of top- k planning when user preferences are not fully known and the model is incomplete.

Example 9. Consider Example 1, in which we identified an optimal plan where the rover first navigates to (7,1) before the drone flies and takes an image at (6,1) and then at (10,1). Finally, the rover equipped with the drone navigates to (0,5). A closer look reveals that the order in which the drone takes the images at the two desired locations is irrelevant to the overall cost of the plan. With top- k planning we can determine both plans, one where the drone flies first to (6,1) and then to (10,1) and vice versa. Based on these alternatives, it can be decided at which of the two locations the images should be taken first, depending on what the user prefers. In addition, there are various plans with different navigation routes of the rover, which allow to choose a different path during the execution of the plan without replanning and thus to react to unforeseen events that were not considered in the model.

7.2 Complexity and Compilability

Clearly, top- k planning is as hard as classical planning, since it is a straightforward generalization. However, the question arises to what extent it is harder to determine multiple plans instead of a single one. Unlike the previously

discussed extension to classical planning, top- k planning is not a direct model extension. The concept of compilability cannot (readily) be applied here, since the output changes from a single plan to a set of plans, which raises a conceptually different decision problem. Thus, Speck, Mattmüller, and Nebel (2020) examined the computational complexity of top- k planning and defined the *bounded top- k plans existence* problem (Definition 16) to answer the question of whether top- k planning is as difficult as classical planning.

Definition 16 (Bounded Top- k Plans Existence). *Bounded top- k plans existence* is the decision problem: Given a planning task Π and two natural numbers ℓ and k , are there at least k different plans of length at most ℓ ?

First of all, Speck, Mattmüller, and Nebel (2020) showed that the bounded top- k plans existence problem in general is PSPACE-complete (Theorem 6), which is surprising because the ordinary bounded plan existence problem, i.e., answering the question whether only one such plan exists, is also PSPACE-complete (Theorem 1).

Theorem 6 (Speck, Mattmüller, and Nebel 2020). *Bounded top- k plans existence* is PSPACE-complete. \square

Considering polynomially bounded plans, it turns out that the bounded top- k plan existence decision problem is PP-hard (Gill 1977; Speck, Mattmüller, and Nebel 2020), while the ordinary bounded plan existence problem is NP-complete (Bylander 1994).

Theorem 7 (Speck, Mattmüller, and Nebel 2020). *Bounded top- k plans existence* is PP-hard if the plan length ℓ is polynomially bounded by the task size, i.e., $\ell \leq p(|\Pi|)$ for some polynomial p . \square

Theorem 7 indicates that top- k planning is much harder in practice than ordinary classical planning. The assumption that both problems have the same complexity when the plan length is polynomially bounded would imply a collapse of the polynomial hierarchy at P^{NP} (Toda 1991), which is considered very unlikely.

7.3 Symbolic Search

Speck, Mattmüller, and Nebel (2020) propose a symbolic search approach to top- k planning that is sound and complete. Recall that symbolic search for classical planning is divided into two phases: 1) a reachability phase, in which states with increasing costs are generated until a goal state is found, and 2) a plan reconstruction phase, in which the search is regressed to reconstruct a corresponding goal path and plan. Since symbolic search expands entire sets of states at once, it is not unusual to find multiple goal states at once. If multiple

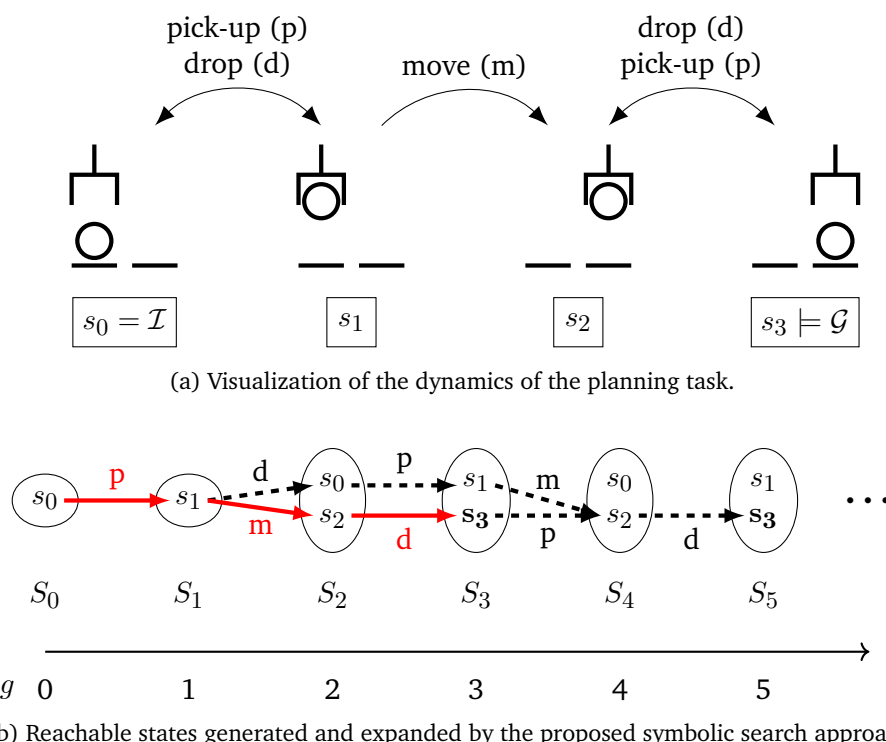


Figure 7.2: Visualization of the GRIPPER planning task and the functioning of the proposed symbolic search approach for top- k planning (Speck, Mattmüller, and Nebel 2020).

goal states are found, then of course multiple plans are found as well. But even if only one goal state is found, it is possible that multiple plans are found that lead to that particular goal state. Usually, symbolic search reconstructs and reports only one such plan, ignoring all others. Based on this observation, Speck, Mattmüller, and Nebel (2020) show that three modifications to ordinary symbolic search lead to a sound and complete symbolic search algorithm for top- k planning. First, they propose that after a goal state is expanded, all plans leading to that goal state are reconstructed. Second, during the search, states are not closed, i.e., all newly generated states with their corresponding reachability costs are insert in the open list without filtering them by already expanded states, otherwise suboptimal plans may be lost. Third and finally, the termination criterion is adjusted so that the algorithm terminates if either k plans are found or the open list contains only states that have already been expanded at least once and are not part of a goal path induced by a plan that has already been found.

Example 10 illustrates the functioning of this modified symbolic search for top- k planning. For simplicity, this example describes symbolic forward search for a task with unit cost.

Example 10 (Speck, Mattmüller, and Nebel 2020). Consider a unit cost planning task with two rooms and a robot with a gripper, as shown in Figure 7.2a. The robot can move from room A to B if it carries the ball, but it cannot return. There is a possibility to pick-up and drop the ball in each room. The goal is to get the ball from room A to room B.

Assuming that the desired number of plans is $k = 3$, Figure 7.2b illustrates the functioning of the proposed symbolic search approach. First, all states S_0 reachable with cost 0, i.e., only the initial state, are expanded resulting in the set of states $S_1 = \{s_1\}$ reachable with cost 1. The subsequent expansion of S_1 yields the state set $S_2 = \{s_0, s_2\}$, whose cost is 2. Note that s_0 is a part of S_2 , although it has been previously expanded and therefore would no longer be considered in ordinary symbolic search. The next expansion yields $S_3 = \{s_1, s_3\}$ with cost 3, leading to the expansion of S_3 , which contains the goal state s_3 . Therefore, the plan reconstruction procedure is executed that yields exactly one plan $\pi_1 = \langle \text{pick-up}, \text{move}, \text{drop} \rangle$ visualized in red in Figure 7.2b. Since a total of 3 plans is desired, the search continues from here until the next goal states are expanded, which occurs with the expansion of S_5 . Plan reconstruction yields two plans with a cost of 5 each, namely $\pi_2 = \langle \text{pick-up}, \text{drop}, \text{pick-up}, \text{move}, \text{drop} \rangle$ and $\pi_3 = \langle \text{pick-up}, \text{move}, \text{drop}, \text{pick-up}, \text{drop} \rangle$. As the desired number of plans is achieved, the algorithm terminates with $\{\pi_1, \pi_2, \pi_3\}$. However, if more plans were desired, the search would continue.

Looking more closely at the reconstruction phase, we can see that we perform an exhaustive greedy backward search using the provided perfect heuristic of the reachability phase. Although this exhaustive search may seem expensive, it is goal-driven due to the perfect heuristic, and each time the initial state is reached, a new plan is created. Finally, Speck, Mattmüller, and Nebel (2020) also define and describe the support of general operator costs, including zero costs, and defined symbolic backward search and symbolic bidirectional search, which can significantly improve planning performance.

Empirical Evaluation. To compare different approaches to top- k planning, we consider k -coverage, i.e., the number of instances for which a planner reports a set of k best plans or reports only $k' < k$ plans but proves that only k' plans exist (see Definition 15). Figure 7.3 compares the k -coverage of the discussed symbolic search approach based on BDDs from Speck, Mattmüller, and Nebel (2020) with other state-of-the-art approaches from the literature on domains of the optimal track of the International Planning Competitions 1998-2018. The x -axis indicates the number of desired plans k , while the y -axis represents the number of instances in which the corresponding k -coverage has been reached. K^* (Aljazzar and Leue 2011; Katz, Sohrabi, Udrea, and Winterer 2018) generates and processes parts of the implicit search tree as needed and can be enhanced by a heuristic. However, the evaluation only includes K^* with




Figure 7.3: The k -coverage (number of instances for which k best plans were found) for explicit and symbolic search algorithms on classical planning domains (Speck, Mattmüller, and Nebel 2020).

the blind heuristic because K^* is very memory intensive, so the performance differences reported by Speck, Mattmüller, and Nebel (2020) between K^* with and without sophisticated heuristics are negligible. Forbid- k (Katz, Sohrabi, Udrea, and Winterer 2018) is an iterative approach that searches for additional plans through a replanning loop that forbids already found plans and preserves all other plans. The underlying search is an orbit space search with structural symmetries (Alkhazraji et al. 2014; Domshlak, Katz, and Shleyfman 2015) and the LM-cut heuristic (Helmert and Domshlak 2009). The three versions of Forbid- k considered differ in a plan reordering step, where different reordering strategies (none, naive, and neighbor) can be used to generate multiple plans from a single plan.

Overall, we can see that symbolic bidirectional search performs as well as Forbid- k when only one plan ($k = 1$) is requested. But already for $k = 2$ symbolic bidirectional search performs best and for $k = 4$ symbolic forward search surpasses Forbid- k , while symbolic bidirectional search already shows a large performance advantage. Most importantly, symbolic search scales much better to larger k than all other approaches. K^* solves only six more instances for $k = 1$ than for $k = 10000$, due to the high memory consumption that is the bottleneck of the approach. The replanning of Forbid- k turns out to be too expensive when a larger number of plans is desired, resulting in a large

performance drop. Finally, Speck, Mattmüller, and Nebel (2020) expect the dominant performance of symbolic search to hold even for very large k until the high number of plan reports is the limiting factor, i.e., the limiting factor is writing the plans to disk.

Discussion

Be happy, but never satisfied.

— Bruce Lee

As we discussed in Section 3.3, symbolic blind search, i.e., without any heuristics, is the dominant search strategy for symbolic search, on par with explicit heuristic search in optimal planning. For this reason, symbolic search provides a good basis for supporting several extensions of classical planning. In this chapter, we discuss the combination of the previously considered extensions and how symbolic search with the introduced modifications can support this setting. Finally, possible future work related to this thesis is discussed.

8.1 Combination of Extensions

As we have seen in the previous chapters, it is possible to extend symbolic search to support various extensions of classical planning. A natural question that arises is, of course, about combining these extensions, which makes intuitive sense since they all capture and extend different aspects of classical planning. The key question is whether the different enhancements to symbolic search can be combined to support several of these planning extensions simultaneously. The short answer to this question is: Yes, with some minor restrictions on the search directions. This is remarkable because the support of only one of these features is quite rare, and the support of more than one feature is almost inexistent in the literature on cost-optimal planning. The reason is the same why very few planning techniques and planners support a single feature, namely that almost all optimal planners are based on heuristic search and it is very difficult to design admissible, informative, and fast to evaluate heuristics that take into account the different extensions to classical planning. We will now briefly discuss the combination of the discussed extensions for classical

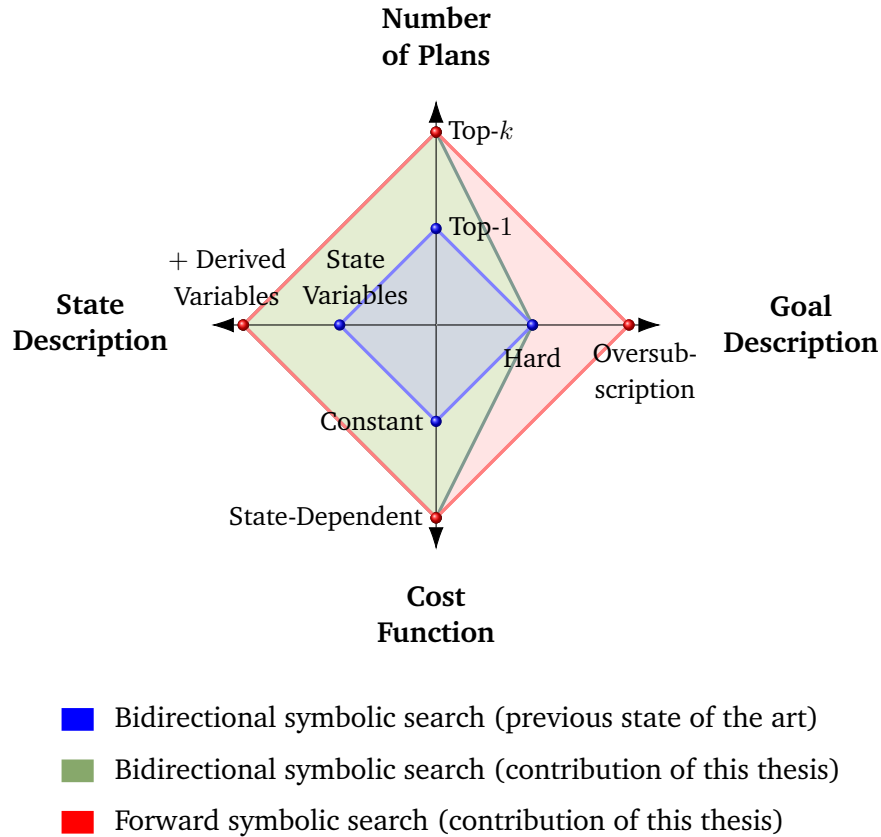


Figure 8.1: Overview of the contribution of this thesis in terms of extending symbolic search to support different expressive extensions of classical planning.

planning and how and why the proposed modifications to symbolic search can support all of these extensions simultaneously.

Model Extensions. Let us first focus on the three discussed model extensions that result in *oversubscription planning with axioms and state-dependent action costs*. This planning formalism is simply defined by combining all the features of the extension and allowing the state-dependent cost function and the utility function to take into account derived variables \mathcal{D} such that the cost function of an operator o is given by $\mathcal{C}_o : \mathcal{S}_E \rightarrow \mathbb{N}_0$ and the utility function $\mathcal{U} : \mathcal{S}_E \rightarrow \mathbb{N}_0$. This makes it possible to concisely encode complex state-dependent action cost functions and state utility functions. In the empirical evaluation of Chapter 5, we have already mentioned that derived variables can be used in a state-dependent action cost function in the Power Supply Restoration (PSR) domain (Hoffmann, Edelkamp, et al. 2006; Thiébaux and Cordier 2001, 2013), to encode whether certain parts of the power system are energized. A closer look,

however, reveals an interesting connection between these model extensions. While derived variables in the cost or utility function can lead to a simple and concise representation, it might be possible to simulate the evaluation of the values of the derived variables by the cost or utility function, which could make the consideration of derived variables in the cost or utility function obsolete. Moreover, it might be possible to simplify a complex cost or utility function by introducing new derived variables and thus shifting the complexity into the computation of the derived variables. This shows that these extensions are closely related and a more detailed analysis of the connection is an exciting question for future work (Section 8.2).

All proposed symbolic encodings of axioms and derived variables in the context of symbolic search (Chapter 4) allow derived variables to exist in the domain of operator cost functions and utility functions. Thus, the proposed enhancements to symbolic search can be directly used for oversubscription planning with axioms and state-dependent action costs. This is especially true for the dominant axiom encoding strategy, symbolic translation, in which each derived variable is replaced by its primary representation (Definition 9). However, as Figure 8.1 illustrates, in oversubscription planning tasks, i.e., in the presence of utilities, only forward symbolic search can be performed, since it is not clear how to perform backward or bidirectional symbolic search efficiently.

Top- k Planning with Model Extensions. The search for the best k plans is independent of the support of axioms and state-dependent action costs and can be easily combined. However, if one combines oversubscription planning and top- k planning, the question arises which are the best k plans. Although this question has never been answered in the literature, it is natural to rank plans first by utility and then by cost, i.e., a plan π is better than a plan π' iff 1) $\mathcal{U}(\pi) > \mathcal{U}(\pi')$ or 2) $\mathcal{U}(\pi) = \mathcal{U}(\pi')$ and $\text{cost}(\pi) < \text{cost}(\pi')$. And indeed, this can be supported with minor modifications to the proposed symbolic search approach by maintaining not only the states with the best utility found so far, but all relevant ones and executing exhaustive plan extraction as done in symbolic search for top- k planning (Figure 8.1).

Overall, the proposed extension of symbolic search yields an optimal, sound, and complete approach to *top- k oversubscription planning with axioms and state-dependent action costs*. As Figure 8.1 illustrates, this not only enhances the applicability and state of the art of symbolic search for classical planning with expressive extensions, but also provides the first planner that can support several of these extensions at once.⁷ Finally, the proposed planning algorithm that supports the interaction of all these features does not suffer from its generality, because if certain features are not present, the approach becomes standard symbolic search.

⁷Available online: <https://github.com/speckdavid/symk>

8.2 Future Work

While we have highlighted some future challenges and research in the individual chapters, here we focus on future research directions related to symbolic (heuristic) search and planning with expressive extensions in general.

We have seen in this thesis different types of decision diagrams such as EVMDDs, ADDs and BDDs used in the context of symbolic search for classical planning. A more detailed comparison in theory and practice could provide more insight into when to use which decision diagrams. Other types of decision diagrams such as Functional Decision Diagrams (Kebschull, Schubert, and Rosenstiel 1992) or Kronecker Functional Decision Diagrams (Drechsler and Becker 1998b) could also lead to an even more concise and efficient representation.

Throughout this thesis, we have assumed a fixed static variable order. Since the order of the variables plays an important role in the efficiency of symbolic search, as the size of the decision diagrams strongly depends on the chosen order of the variables, it is important to find good orders. While it is known that the computation of an optimal order for decision diagrams is co-NP-complete (Bryant 1986), it is still a challenging question whether and how to find good orders in practice (Kissmann and Hoffmann 2013, 2014). Moreover, dynamic reordering techniques (Rudell 1993) are used in other fields, such as model checking (Yang et al. 1998) or logic synthesis (Scholl et al. 1999), which have been explored up to now only sparsely in planning (Kissmann, Edelkamp, and Hoffmann 2014; Kissmann and Hoffmann 2014).

The efficient integration of heuristics into symbolic search is still an open research task. We believe that the main goal should be to keep the decision diagrams involved as small as possible. Fišer, Torralba, and Hoffmann (2021) introduced operator potentials based on potential heuristics (Pommerening et al. 2015), which seem to provide concise representation in symbolic forward search, while efficiently pruning states based on their heuristic values. More work certainly needs to be done here to better understand the search behavior of symbolic heuristic search. In particular, it is interesting to see if it is possible to find heuristics that give any size guarantees for the decision diagrams involved, and if and how the results can be generalized to symbolic bidirectional heuristic search.

Considering the discussed extensions of classical planning, an in-depth theoretical and empirical evaluation of the combination of several such extensions and the symbolic search approach presented for this purpose is certainly a promising research direction. In particular, it is interesting to examine the relationship between different extensions. It might be possible that one extension can be compiled away (with small overhead) in the presence of another extension. Moreover, several of these extensions are not yet covered by more expressive planning formalisms, such as fully observable nondeterministic planning (Cimatti et al. 2003), partially observable nondeterministic planning

(Bertoli et al. 2006; Speck, Ortlieb, and Mattmüller 2015), or hierarchical task network planning (Erol, Hendler, and Nau 1996; Geier and Bercher 2011). An analysis of how and to what extent all these features can enhance these planning formalisms is an important future line of research, and whether the proposed ideas can also be used for symbolic search approaches for these settings (Behnke and Speck 2021; Bertoli et al. 2006; Kissmann and Edelkamp 2009).

Conclusion

Perhaps the best test of [an agent's] intelligence is [its] capacity for making a summary.

— Lytton Strachey

We theoretically and empirically evaluated symbolic search for cost-optimal planning with expressive extensions. First, we analyzed the search behavior of symbolic heuristic search in form of BDDA^{*} and revealed an unknown or at least often overlooked fundamental problem: We showed that using a heuristic does not always improve the search performance of BDDA^{*}. In general, even the perfect heuristic can deteriorate search performance exponentially. Second, we analyzed classical planning with different expressive extensions, such as planning with axioms and derived variables, planning with state-dependent action costs, oversubscription planning, and top- k planning. The discussed and presented results on the computational complexity and compilability of these planning formalisms show that native support for the relevant planning extensions is needed to solve many real-world problems. Based on this observation, we proposed symbolic blind search to solve such planning formalisms with expressive extensions. More specifically, we have proposed symbolic search approaches that provide optimal, sound, and complete algorithms for these planning formalisms. Our empirical evaluations show that the presented symbolic search approaches perform favorably in all these planning settings compared with other state-of-the-art approaches. Finally, we analyzed the combination of all these classical planning extensions, i.e., top- k oversubscription planning with axioms and state-dependent action costs, and show how symbolic search with the introduced modifications can support this planning formalism.

List of Figures

1.1	Overview of extension for classical planning.	2
2.1	Visualization of the Mars rover planning task (running example).	9
2.2	Visualization of different decision diagrams.	14
3.1	Illustration of BDDA*.	16
3.2	Visualization of two BDDs with exponential size difference.	18
3.3	Comparison of BDDA* with two different heuristics.	20
4.1	Overview of extension for classical planning (axioms).	24
4.2	\mathcal{O} -based encoding of axioms.	27
4.3	\mathcal{V} -based encoding of axioms.	28
4.4	Symbolic translation encoding of axioms.	29
5.1	Overview of extension for classical planning (cost function).	32
5.2	Visualization of symbolic representations of an operator with state-dependent action costs.	37
6.1	Overview of extension for classical planning (goal description).	40
6.2	Visualization of an oversubscription planning task.	44
6.3	Coverage over time for oversubscription planning.	46
7.1	Overview of extension for classical planning (top- k).	48
7.2	Visualization of symbolic search for top- k planning.	51
7.3	The k -coverage for top- k planning.	53
8.1	Overview of the contribution of this thesis.	56

List of Tables

4.1	Coverage for planning with axioms.	30
5.1	Compilability results for planning with state-dependent action cost.	35
5.2	Coverage for planning with state-dependent action costs.	38
6.1	Coverage for oversubscription planning.	45

Bibliography

- Aghighi, Meysam and Christer Bäckström (2015). “Cost-Optimal and Net-Benefit Planning - A Parameterised Complexity View”. In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*. Ed. by Qiang Yang and Michael Wooldridge. AAAI Press, pp. 1487–1493.
- Aghighi, Meysam and Peter Jonsson (2014). “Oversubscription Planning: Complexity and Compilability”. In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*. Ed. by Carla E. Brodley and Peter Stone. AAAI Press, pp. 2221–2227.
- Aljazzar, Husain and Stefan Leue (2011). “K*: A Heuristic Search Algorithm for Finding the K Shortest Paths”. In: *Artificial Intelligence* 175.18, pp. 2129–2154.
- Alkhazraji, Yusra, Michael Katz, Robert Mattmüller, Florian Pommerening, Alexander Shleyfman, and Martin Wehrle (2014). “Metis: Arming Fast Downward with Pruning and Incremental Computation”. In: *Eighth International Planning Competition (IPC-8): Planner Abstracts*, pp. 88–92.
- Apt, Krzysztof R., Howard A. Blair, and Adrian Walker (1988). “Towards a Theory of Declarative Knowledge”. In: *Foundations of Deductive Databases and Logic Programming*. Ed. by Jack Minker. Morgan Kaufmann, pp. 89–148.
- Bäckström, Christer and Bernhard Nebel (1995). “Complexity Results for SAS⁺ Planning”. In: *Computational Intelligence* 11.4, pp. 625–655.
- Bahar, R. Iris, Erica A. Frohm, Charles M. Gaona, Gary D. Hachtel, Enrico Macii, Abelardo Pardo, and Fabio Somenzi (1997). “Algebraic Decision Diagrams and Their Applications”. In: *Formal Methods in System Design* 10.2–3, pp. 171–206.

- Becker, Bernd and Paul Molitor (2008). *Technische Informatik - eine einführende Darstellung*. Oldenbourg Wissenschaftsverlag.
- Behnke, Gregor and David Speck (2021). “Symbolic Search for Total-Order HTN Planning”. In: *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*. Ed. by Kevin Leyton-Brown and Mausam. AAAI Press, pp. 11744–11754.
- Bertoli, Piergiorgio, Alessandro Cimatti, Marco Roveri, and Paolo Traverso (2006). “Strong Planning under Partial Observability”. In: *Artificial Intelligence* 170, pp. 337–384.
- Bock, Frederick, Harold Kanter, and John Haynes (1957). *An Algorithm (the r-th Best Path Algorithm) for Finding and Ranking Paths through a Network*. Armour Research Foundation.
- Bonet, Blai and Hector Geffner (1999). “Planning as Heuristic Search: New Results”. In: *Recent Advances in AI Planning. 5th European Conference on Planning (ECP 1999)*. Ed. by Susanne Biundo and Maria Fox. Vol. 1809. Lecture Notes in Artificial Intelligence. Heidelberg: Springer, pp. 360–372.
- Brace, Karl S., Richard Rudell, and Randal E. Bryant (1990). “Efficient Implementation of a BDD Package”. In: *Proceedings of the 27th ACM/IEEE Design Automation Conference (DAC 1990)*. Ed. by Richard C. Smith, pp. 40–45.
- Bryant, Randal E. (1985). “Symbolic Manipulation of Boolean Functions Using a Graphical Representation”. In: *Proceedings of the 22nd ACM/IEEE Conference on Design Automation (DAC 1985)*. Ed. by Hillel Ofek and Lawrence A. O’Neill, pp. 688–694.
- Bryant, Randal E. (1986). “Graph-Based Algorithms for Boolean Function Manipulation”. In: *IEEE Transactions on Computers* 35.8, pp. 677–691.
- Burch, Jerry R., Edmund M. Clarke, David E. Long, Kenneth L. McMillan, and David L. Dill (1994). “Symbolic Model Checking for Sequential Circuit Verification”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 13.4, pp. 401–424.
- Bylander, Tom (1994). “The Computational Complexity of Propositional STRIPS Planning”. In: *Artificial Intelligence* 69.1–2, pp. 165–204.
- Bylander, Tom (1997). “A Linear Programming Heuristic for Optimal Planning”. In: *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI 1997)*. AAAI Press, pp. 694–699.
- Ciardo, Gianfranco and Radu Siminiceanu (2002). “Using Edge-Valued Decision Diagrams for Symbolic Generation of Shortest Paths”. In: *Proceedings of the Fourth International Conference on Formal Methods in Computer-Aided*

- Design (FMCAD 2002)*. Ed. by Mark Aagaard and John W. O’Leary. Springer, pp. 256–273.
- Cimatti, Alessandro, Marco Pistore, Marco Roveri, and Paolo Traverso (2003). “Weak, Strong, and Strong Cyclic Planning via Symbolic Model Checking”. In: *Artificial Intelligence* 147, pp. 35–84.
- Corraya, Sumitra, Florian Geißer, David Speck, and Robert Mattmüller (2019). “An Empirical Study of the Usefulness of State-Dependent Action Costs in Planning”. In: *Proceedings of the 42nd Annual German Conference on Artificial Intelligence (KI 2019)*. Ed. by Christoph Benzmüller and Heiner Stuckenschmidt. Lecture Notes in Artificial Intelligence. Springer, pp. 123–130.
- Cui, Hao and Roni Khardon (2018). “The SOGBOFA System in IPC 2018: Lifted BP for Conformant Approximation of Stochastic Planning”. In: *Sixth International Probabilistic Planning Competition (IPPC-6): Planner Abstracts*, pp. 1–6.
- Dijkstra, Edsger W. (1959). “A Note on Two Problems in Connexion with Graphs”. In: *Numerische Mathematik* 1, pp. 269–271.
- Domshlak, Carmel, Michael Katz, and Alexander Shleyfman (2015). *Symmetry Breaking in Deterministic Planning as Forward Search: Orbit Space Search Algorithm*. Tech. rep. IS/IE-2015-03. Technion.
- Domshlak, Carmel and Vitaly Mirkis (2015). “Deterministic Oversubscription Planning as Heuristic Search: Abstractions and Reformulations”. In: *Journal of Artificial Intelligence Research* 52, pp. 97–169.
- Drechsler, Rolf and Bernd Becker (1998a). *Binary Decision Diagrams - Theory and Implementation*. Springer.
- Drechsler, Rolf and Bernd Becker (1998b). “Ordered Kronecker Functional Decision Diagrams – a Data Structure for Representation and Manipulation of Boolean Functions”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 17.10, pp. 965–973.
- Drexler, Dominik, Jendrik Seipp, and David Speck (2021). “Subset-Saturated Transition Cost Partitioning”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 131–139.
- Drexler, Dominik, David Speck, and Robert Mattmüller (2020). “Subset-Saturated Transition Cost Partitioning for Optimal Classical Planning”. In: *ICAPS 2020 Workshop on Heuristics and Search for Domain-independent Planning (HSDIP)*, pp. 23–31.

- Edelkamp, Stefan (2002). “Symbolic Pattern Databases in Heuristic Search Planning”. In: *Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling (AIPS 2002)*. Ed. by Malik Ghallab, Joachim Hertzberg, and Paolo Traverso. AAAI Press, pp. 274–283.
- Edelkamp, Stefan (2003). “Limits and Possibilities of PDDL for Model Checking Software”. In: *Proceedings of the ICAPS 2003 Workshop on the Competition: Impact, Organisation, Evaluation, Benchmarks*. Ed. by Stefan Edelkamp and Jörg Hoffmann.
- Edelkamp, Stefan and Peter Kissmann (2009). “Optimal Symbolic Planning with Action Costs and Preferences”. In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*. Ed. by Craig Boutilier. AAAI Press, pp. 1690–1695.
- Edelkamp, Stefan and Peter Kissmann (2011). “On the Complexity of BDDs for State Space Search: A Case Study in Connect Four”. In: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*. Ed. by Wolfram Burgard and Dan Roth. AAAI Press, pp. 18–23.
- Edelkamp, Stefan, Peter Kissmann, and Álvaro Torralba (2015). “BDDs Strike Back (in AI Planning)”. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*. Ed. by Blai Bonet and Sven Koenig. AAAI Press, pp. 4320–4321.
- Edelkamp, Stefan and Frank Reffel (1998). “OBDDs in Heuristic Search”. In: *Proceedings of the 22nd Annual German Conference on Artificial Intelligence (KI 1998)*. Ed. by Otthein Herzog and Andreas Günter. Vol. 1504. Lecture Notes in Computer Science. Springer, pp. 81–92.
- Eifler, Rebecca, Marcel Steinmetz, Álvaro Torralba, and Jörg Hoffmann (2020). “Plan-Space Explanation via Plan-Property Dependencies: Faster Algorithms & More Powerful Properties”. In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*. IJCAI, pp. 4091–4097.
- Erol, Kutluhan, James A. Hendler, and Dana S. Nau (1996). “Complexity Results for HTN Planning”. In: *Annals of Mathematics and Artificial Intelligence (AMAI)* 18.1, pp. 69–93.
- Fikes, Richard E. and Nils J. Nilsson (1971). “STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving”. In: *Artificial Intelligence* 2, pp. 189–208.
- Fišer, Daniel, Álvaro Torralba, and Jörg Hoffmann (2021). “Introducing Operator-Potential Heuristics for Symbolic Search”. In: *ICAPS 2021 Workshop on Heuristics and Search for Domain-independent Planning (HSDIP)*.

- Franco, Santiago, Levi H. S. Lelis, and Mike Barley (2018). “The Complementary2 Planner in the IPC 2018”. In: *Ninth International Planning Competition (IPC-9): Planner Abstracts*, pp. 32–36.
- Franco, Santiago, Álvaro Torralba, Levi H. S. Lelis, and Mike Barley (2017). “On Creating Complementary Pattern Databases”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*. Ed. by Carles Sierra. IJCAI, pp. 4302–4309.
- García Olaya, Angel, Tomás de la Rosa, and Daniel Borrajo (2021). “Selecting Goals in Oversubscription Planning Using Relaxed Plans”. In: *Artificial Intelligence* 291, p. 103414.
- Gazen, B. Cenk and Craig A. Knoblock (1997). “Combining the Expressivity of UCPOP with the Efficiency of Graphplan”. In: *Recent Advances in AI Planning. 4th European Conference on Planning (ECP 1997)*. Ed. by Sam Steel and Rachid Alami. Vol. 1348. Lecture Notes in Artificial Intelligence. Springer, pp. 221–233.
- Geier, Thomas and Pascal Bercher (2011). “On the Decidability of HTN Planning with Task Insertion”. In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*. Ed. by Toby Walsh. AAAI Press, pp. 1955–1961.
- Geißer, Florian (2018). “On Planning with State-Dependent Action Costs”. PhD thesis. University of Freiburg.
- Geißer, Florian, Thomas Keller, and Robert Mattmüller (2015). “Delete Relaxations for Planning with State-Dependent Action Costs”. In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*. Ed. by Qiang Yang and Michael Wooldridge. AAAI Press, pp. 1573–1579.
- Geißer, Florian, Thomas Keller, and Robert Mattmüller (2016). “Abstractions for Planning with State-Dependent Action Costs”. In: *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling (ICAPS 2016)*. Ed. by Amanda Coles, Andrew Coles, Stefan Edelkamp, Daniele Magazzeni, and Scott Sanner. AAAI Press, pp. 140–148.
- Geißer, Florian and David Speck (2018). “Prost-DD – Utilizing Symbolic Classical Planning in THTS”. In: *Sixth International Probabilistic Planning Competition (IPPC-6): Planner Abstracts*, pp. 13–16.
- Geißer, Florian, David Speck, and Thomas Keller (2019). “An Analysis of the Probabilistic Track of the IPC 2018”. In: *ICAPS 2019 Workshop on the International Planning Competition (WIPC)*, pp. 27–35.
- Geißer, Florian, David Speck, and Thomas Keller (2020). “Trial-Based Heuristic Tree Search for MDPs with Factored Action Spaces”. In: *Proceedings of*

- the 13th Annual Symposium on Combinatorial Search (SoCS 2020)*. Ed. by Daniel Harabor and Mauro Vallati. AAAI Press, pp. 38–47.
- Ghosh, Kamalesh, Pallab Dasgupta, and S. Ramesh (2015). “Automated Planning as an Early Verification Tool for Distributed Control”. In: *Journal of Automated Reasoning* 54.1, pp. 31–68.
- Gill, John (1977). “Computational Complexity of Probabilistic Turing Machines”. In: *SIAM Journal on Computing* 6.4, pp. 675–695.
- Hansen, Eric A., Rong Zhou, and Zhengzhu Feng (2002). “Symbolic Heuristic Search Using Decision Diagrams”. In: *Proceedings of the 5th International Symposium on Abstraction, Reformulation and Approximation (SARA 2002)*. Ed. by Sven Koenig and Robert C. Holte. Vol. 2371. Lecture Notes in Artificial Intelligence. Springer, pp. 83–98.
- Hart, Peter E., Nils J. Nilsson, and Bertram Raphael (1968). “A Formal Basis for the Heuristic Determination of Minimum Cost Paths”. In: *IEEE Transactions on Systems Science and Cybernetics* 4.2, pp. 100–107.
- Haslum, Patrik, Franc Ivankovic, Miquel Ramirez, Dan Gordon, Sylvie Thiébaux, Vikas Shivashankar, and Dana S. Nau (2018). “Extending Classical Planning with State Constraints: Heuristics and Search for Optimal Planning”. In: *Journal of Artificial Intelligence Research* 62, pp. 373–431.
- Helmert, Malte (2006). “The Fast Downward Planning System”. In: *Journal of Artificial Intelligence Research* 26, pp. 191–246.
- Helmert, Malte (2008). *Understanding Planning Tasks – Domain Complexity and Heuristic Decomposition*. Vol. 4929. Lecture Notes in Artificial Intelligence. Springer.
- Helmert, Malte and Carmel Domshlak (2009). “Landmarks, Critical Paths and Abstractions: What’s the Difference Anyway?” In: *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS 2009)*. Ed. by Alfonso Gerevini, Adele Howe, Amedeo Cesta, and Ioannis Refanidis. AAAI Press, pp. 162–169.
- Helmert, Malte and Hauke Lasinger (2010). “The Scanalyzer Domain: Greenhouse Logistics as a Planning Problem”. In: *Proceedings of the Twentieth International Conference on Automated Planning and Scheduling (ICAPS 2010)*. Ed. by Ronen Brafman, Hector Geffner, Jörg Hoffmann, and Henry Kautz. AAAI Press, pp. 234–237.
- Hoffmann, Jörg and Stefan Edelkamp (2005). “The Deterministic Part of IPC-4: An Overview”. In: *Journal of Artificial Intelligence Research* 24, pp. 519–579.
- Hoffmann, Jörg, Stefan Edelkamp, Sylvie Thiébaux, Roman Englert, Frederico dos Santos Liporace, and Sebastian Trüg (2006). “Engineering Benchmarks

- for Planning: the Domains Used in the Deterministic Part of IPC-4". In: *Journal of Artificial Intelligence Research* 26, pp. 453–541.
- Höller, Daniel (2021). "Translating Totally Ordered HTN Planning Problems to Classical Planning Problems Using Regular Approximation of Context-Free Languages". In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 159–167.
- Ivankovic, Franc, Dan Gordon, and Patrik Haslum (2019). "Planning with Global State Constraints and State-Dependent Action Costs". In: *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*. Ed. by Nir Lipovetzky, Eva Onaindia, and David E. Smith. AAAI Press, pp. 232–236.
- Ivankovic, Franc and Patrik Haslum (2015). "Optimal Planning with Axioms". In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*. Ed. by Qiang Yang and Michael Wooldridge. AAAI Press, pp. 1580–1586.
- Jensen, Rune M., Manuela M. Veloso, and Randal E. Bryant (2008). "State-Set Branching: Leveraging BDDs for Heuristic Search". In: *Artificial Intelligence* 172.2–3, pp. 103–139.
- Karpas, Erez and Daniele Magazzeni (2020). "Automated Planning for Robotics". In: *Annual Review of Control, Robotics, and Autonomous Systems* 3, pp. 417–439.
- Katz, Michael and Emil Keyder (2019). "A* Search and Bound-Sensitive Heuristics for Oversubscription Planning". In: *ICAPS 2019 Workshop on Heuristics and Search for Domain-independent Planning (HSDIP)*.
- Katz, Michael, Emil Keyder, Florian Pommerening, and Dominik Winterer (2019a). "Oversubscription Planning as Classical Planning with Multiple Cost Functions". In: *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*. Ed. by Nir Lipovetzky, Eva Onaindia, and David E. Smith. AAAI Press, pp. 237–245.
- Katz, Michael, Emil Keyder, Florian Pommerening, and Dominik Winterer (2019b). *PDDL Benchmarks for Oversubscription Planning*. <https://doi.org/10.5281/zenodo.2576024>.
- Katz, Michael and Vitaly Mirkis (2016). "In Search of Tractability for Partial Satisfaction Planning". In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*. Ed. by Subbarao Kambhampati. AAAI Press, pp. 3154–3160.
- Katz, Michael, Shirin Sohrabi, and Octavian Udrea (2020). "Top-Quality Planning: Finding Practically Useful Sets of Best Plans". In: *Proceedings of the*

- Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*. Ed. by Vincent Conitzer and Fei Sha. AAAI Press, pp. 9900–9907.
- Katz, Michael, Shirin Sohrabi, Octavian Udrea, and Dominik Winterer (2018). “A Novel Iterative Approach to Top-k Planning”. In: *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*. Ed. by Mathijs de Weerd, Sven Koenig, Gabriele Röger, and Matthijs Spaan. AAAI Press, pp. 132–140.
- Katz, Michael, Kavitha Srinivas, Shirin Sohrabi, Mark Feblowitz, Octavian Udrea, and Oktie Hassanzadeh (2021). “Scenario Planning in the Wild: A Neuro-Symbolic Approach”. In: *ICAPS 2021 Workshop on Planning for Financial Services (FinPlan)*, pp. 403–411.
- Kebschull, Udo, Endric Schubert, and Wolfgang Rosenstiel (1992). “Multilevel Logic Synthesis Based on Functional Decision Diagrams”. In: *Proceedings of the 29th ACM/IEEE Design Automation Conference (DAC 1992)*. Ed. by Daniel G. Schweikert, pp. 43–47.
- Keller, Thomas and Patrick Eyerich (2012). “PROST: Probabilistic Planning Based on UCT”. In: *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling (ICAPS 2012)*. Ed. by Lee McCluskey, Brian Williams, José Reinaldo Silva, and Blai Bonet. AAAI Press, pp. 119–127.
- Keller, Thomas and Florian Geißer (2015). “Better Be Lucky Than Good: Exceeding Expectations in MDP Evaluation”. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*. Ed. by Blai Bonet and Sven Koenig. AAAI Press, pp. 3540–3547.
- Keller, Thomas, Florian Pommerening, Jendrik Seipp, Florian Geißer, and Robert Mattmüller (2016). “State-Dependent Cost Partitionings for Cartesian Abstractions in Classical Planning”. In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*. Ed. by Subbarao Kambhampati. AAAI Press, pp. 3161–3169.
- Keyder, Emil and Hector Geffner (2009). “Soft Goals Can Be Compiled Away”. In: *Journal of Artificial Intelligence Research* 36, pp. 547–556.
- Kissmann, Peter (2012). “Symbolic Search in Planning and General Game Playing”. PhD thesis. University of Bremen.
- Kissmann, Peter and Stefan Edelkamp (2009). “Solving Fully-Observable Non-Deterministic Planning Problems via Translation into a General Game”. In: *Proceedings of the 32nd Annual German Conference on Artificial Intelligence (KI 2009)*. Ed. by Bärbel Mertsching, Marcus Hund, and Zaheer Aziz. Vol. 5803. Lecture Notes in Artificial Intelligence. Springer, pp. 1–8.

- Kissmann, Peter and Stefan Edelkamp (2011). “Improving Cost-Optimal Domain-Independent Symbolic Planning”. In: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*. Ed. by Wolfram Burgard and Dan Roth. AAAI Press, pp. 992–997.
- Kissmann, Peter, Stefan Edelkamp, and Jörg Hoffmann (2014). “Gamer and Dynamic-Gamer – Symbolic Search at IPC 2014”. In: *Eighth International Planning Competition (IPC-8): Planner Abstracts*, pp. 77–84.
- Kissmann, Peter and Jörg Hoffmann (2013). “What’s in It for My BDD? On Causal Graphs and Variable Orders in Planning”. In: *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*. Ed. by Daniel Borrajo, Subbarao Kambhampati, Angelo Oddi, and Simone Fratini. AAAI Press, pp. 327–331.
- Kissmann, Peter and Jörg Hoffmann (2014). “BDD Ordering Heuristics for Classical Planning”. In: *Journal of Artificial Intelligence Research* 51, pp. 779–804.
- Koehler, Jana and Kilian Schuster (2000). “Elevator Control as a Planning Problem”. In: *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS 2000)*. Ed. by Steve Chien, Subbarao Kambhampati, and Craig A. Knoblock. AAAI Press, pp. 331–338.
- Koller, Alexander and Jörg Hoffmann (2010). “Waking Up a Sleeping Rabbit: On Natural-Language Sentence Generation with FF”. In: *Proceedings of the Twentieth International Conference on Automated Planning and Scheduling (ICAPS 2010)*. Ed. by Ronen Brafman, Hector Geffner, Jörg Hoffmann, and Henry Kautz. AAAI Press, pp. 238–241.
- Kominis, Filippos and Hector Geffner (2015). “Beliefs in Multiagent Planning: From One Agent to Many”. In: *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling (ICAPS 2015)*. Ed. by Ronen Brafman, Carmel Domshlak, Patrik Haslum, and Shlomo Zilberstein. AAAI Press, pp. 147–155.
- Lai, Yung-Te, Massoud Pedram, and Sarma B. K. Vrudhula (1996). “Formal Verification Using Edge-Valued Binary Decision Diagrams”. In: *IEEE Transactions on Computers* 45.2, pp. 247–255.
- LaValle, Steven M. (2006). *Planning Algorithms*. Cambridge University Press.
- Lozano-Pérez and Leslie P. Kaelbling (2014). “A Constraint-Based Method for Solving Sequential Manipulation Planning Problems”. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014)*, pp. 3684–3691.
- Mattmüller, Robert, Florian Geißer, Benedict Wright, and Bernhard Nebel (2018). “On the Relationship Between State-Dependent Action Costs and

- Conditional Effects in Planning”. In: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*. AAAI Press, pp. 6237–6245.
- McDermott, Drew, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins (1998). *PDDL – The Planning Domain Definition Language – Version 1.2*. Tech. rep. CVC TR-98-003/DCS TR-1165. Yale University: Yale Center for Computational Vision and Control.
- McMillan, Kenneth L. (1993). *Symbolic Model Checking*. Kluwer Academic Publishers.
- Mirkis, Vitaly and Carmel Domshlak (2013). “Abstractions for Oversubscription Planning”. In: *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*. Ed. by Daniel Borrajo, Subbarao Kambhampati, Angelo Oddi, and Simone Fratini. AAAI Press, pp. 153–161.
- Miura, Shuwa and Alex Fukunaga (2017). “Automatically Extracting Axioms in Classical Planning”. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI 2017)*. Ed. by Satinder Singh and Shaul Markovitch. AAAI Press, pp. 4973–4974.
- Moraru, Ionut, Stefan Edelkamp, Santiago Franco, and Moises Martinez (2019). “Simplifying Automated Pattern Selection for Planning with Symbolic Pattern Databases”. In: *Proceedings of the 42nd Annual German Conference on Artificial Intelligence (KI 2019)*. Ed. by Christoph Benzmüller and Heiner Stuckenschmidt. Lecture Notes in Artificial Intelligence. Springer, pp. 249–263.
- Muller, Daniel and Erez Karpas (2018). “Value Driven Landmarks for Oversubscription Planning”. In: *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*. Ed. by Mathijs de Weerd, Sven Koenig, Gabriele Röger, and Matthijs Spaan. AAAI Press, pp. 171–179.
- Nebel, Bernhard (2000). “On the Compilability and Expressive Power of Propositional Planning Formalisms”. In: *Journal of Artificial Intelligence Research* 12, pp. 271–315.
- Newell, Allen and Herbert A. Simon (1963). “GPS: A Program that Simulates Human Thought”. In: *Computers and Thought*. Ed. by E. A. Feigenbaum and J. Feldman. Oldenbourg, pp. 279–293.
- Nguyen, Tuan Anh, Minh Binh Do, Alfonso Gerevini, Ivan Serina, Biplav Srivastava, and Subbarao Kambhampati (2012). “Generating Diverse Plans

- to Handle Unknown and Partially Known User Preferences”. In: *Artificial Intelligence* 190, pp. 1–31.
- Nilsson, Nils J. (1984). *Shakey the Robot*. Tech. rep. 323. AI Center, SRI International. Menlo Park, CA, USA.
- Pearl, Judea (1984). *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley.
- Pohl, Ira (1969). *Bi-Directional And Heuristic Search In Path Problems*. Tech. rep. 104. Stanford Linear Accelerator Center.
- Pommerening, Florian, Malte Helmert, Gabriele Röger, and Jendrik Seipp (2015). “From Non-Negative to General Operator Cost Partitioning”. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*. Ed. by Blai Bonet and Sven Koenig. AAAI Press, pp. 3335–3341.
- Puterman, Martin L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- Reif, John H. (1979). “Error Detection with Directed Symbolic Model Checking”. In: *Proceedings of the 20th Annual Symposium on Foundations of Computer Science (FOCS 1979)*, pp. 421–427.
- Ren, Tianyu, Georgia Chalvatzaki, and Jan Peters (2021). *Extended Task and Motion Planning of Long-Horizon Robot Manipulation*. arXiv:2103.05456 [cs.RO].
- Riabov, Anton V., Shirin Sohrabi, and Octavian Udrea (2014). “New Algorithms for the Top-k Planning Problem”. In: *ICAPS 2014 Scheduling and Planning Applications woRKshop*, pp. 10–16.
- Rudell, Richard (1993). “Dynamic Variable Ordering for Ordered Binary Decision Diagrams”. In: *Proceedings of the 1993 IEEE/ACM International Conference on Computer-Aided Design (ICCAD 1993)*. Ed. by Michael R. Lightner and Jochen A. G. Jess, pp. 42–47.
- Russell, Stuart and Peter Norvig (2003). *Artificial Intelligence — A Modern Approach*. Prentice Hall.
- Sanner, Scott (2010). *Relational Dynamic Influence Diagram Language (RDDL): Language Description*.
- Savitch, Walter J. (1970). “Relationships Between Nondeterministic and Deterministic Tape Complexities”. In: *Journal of Computer and System Sciences* 4, pp. 177–192.

- Scholl, Christoph, Dirk Möller, Paul Molitor, and Rolf Drechsler (1999). “BDD Minimization Using Symmetries”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 18.2, pp. 81–100.
- Seimetz, Valentin, Rebecca Eifler, and Jörg Hoffmann (2021). “Learning Temporal Plan Preferences from Examples: An Empirical Study”. In: *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI 2021)*. Ed. by Zhi-Hua Zhou. IJCAI, pp. 4160–4166.
- Sievers, Silvan, Michael Katz, Shirin Sohrabi, Horst Samulowitz, and Patrick Ferber (2019). “Deep Learning for Cost-Optimal Planning: Task-Dependent Planner Selection”. In: *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*. AAAI Press, pp. 7715–7723.
- Siminicéanu, Radu and Pierre Roux (2010). “Model Checking with Edge-Valued Decision Diagrams”. In: *Proceedings of the Second NASA Formal Methods Symposium (NFM 2010)*, pp. 222–226.
- Smith, David E. (2004). “Choosing Objectives in Over-Subscription Planning”. In: *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*. Ed. by Shlomo Zilberstein, Jana Koehler, and Sven Koenig. AAAI Press, pp. 393–401.
- Sohrabi, Shirin, Anton V. Riabov, Michael Katz, and Octavian Udrea (2018). “An AI Planning Solution to Scenario Generation for Enterprise Risk Management”. In: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*. AAAI Press, pp. 160–167.
- Sohrabi, Shirin, Anton V. Riabov, and Octavian Udrea (2016). “Plan Recognition as Planning Revisited”. In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*. Ed. by Subbarao Kambhampati. AAAI Press, pp. 3258–3264.
- Somenzi, Fabio (2015). *CUDD: CU Decision Diagram Package – Release 3.0.0*. <https://github.com/ivmai/cudd>. Accessed: 2020-02-20.
- Speck, David, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (June 5, 2020a). “Device and Method for Planning and Operation of a Technical System”. European pat. req. EP 20178576.3 – 1203. Robert Bosch GmbH. Filed.
- Speck, David, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (2020b). “Learning Heuristic Selection with Dynamic Algorithm Configuration”. In: *ICAPS 2020 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL)*, pp. 61–69.
- Speck, David, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (Apr. 28, 2021a). “Device and Method for Planning and Operation

- of a Technical System”. U.S. pat. req. 17/242,790. Robert Bosch GmbH. Filed.
- Speck, David, André Biedenkapp, Frank Hutter, Robert Mattmüller, and Marius Lindauer (2021b). “Learning Heuristic Selection with Dynamic Algorithm Configuration”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 597–605.
- Speck, David, David Borukhson, Robert Mattmüller, and Bernhard Nebel (2021). “On the Compilability and Expressive Power of State-Dependent Action Costs”. In: *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*. Ed. by Robert P. Goldman, Susanne Biundo, and Michael Katz. AAAI Press, pp. 358–366.
- Speck, David, Christian Dornhege, and Wolfram Burgard (2017). “Shakey 2016 – How Much Does it Take to Redo Shakey the Robot?” In: *IEEE Robotics and Automation Letters (RA-L)* 2.2, pp. 1203–1209.
- Speck, David, Florian Geißer, and Robert Mattmüller (2018a). “Symbolic Planning with Edge-Valued Multi-Valued Decision Diagrams”. In: *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*. Ed. by Mathijs de Weerd, Sven Koenig, Gabriele Röger, and Matthijs Spaan. AAAI Press, pp. 250–258.
- Speck, David, Florian Geißer, and Robert Mattmüller (2018b). “SYMPLE: Symbolic Planning Based on EVMDDs”. In: *Ninth International Planning Competition (IPC-9): Planner Abstracts*, pp. 91–94.
- Speck, David, Florian Geißer, and Robert Mattmüller (2020). “When Perfect Is Not Good Enough: On the Search Behaviour of Symbolic Heuristic Search”. In: *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)*. Ed. by J. Christopher Beck, Erez Karpas, and Shirin Sohrabi. AAAI Press, pp. 263–271.
- Speck, David, Florian Geißer, Robert Mattmüller, and Álvaro Torralba (2019). “Symbolic Planning with Axioms”. In: *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*. Ed. by Nir Lipovetzky, Eva Onaindia, and David E. Smith. AAAI Press, pp. 464–472.
- Speck, David and Michael Katz (2021). “Symbolic Search for Oversubscription Planning”. In: *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*. Ed. by Kevin Leyton-Brown and Mausam. AAAI Press, pp. 11972–11980.
- Speck, David, Robert Mattmüller, and Bernhard Nebel (2020). “Symbolic Top-k Planning”. In: *Proceedings of the Thirty-Fourth AAAI Conference on Artificial*

- Intelligence (AAAI 2020)*. Ed. by Vincent Conitzer and Fei Sha. AAAI Press, pp. 9967–9974.
- Speck, David, Manuela Ortlieb, and Robert Mattmüller (2015). “Necessary Observations in Nondeterministic Planning”. In: *Proceedings of the 38th Annual German Conference on Artificial Intelligence (KI 2015)*. Ed. by Steffen Hölldobler, Markus Krötzsch, Rafael Peñaloza-Nyssen, and Sebastian Rudolph. Vol. 9324. Lecture Notes in Artificial Intelligence. Springer, pp. 181–193.
- Thiébaux, Sylvie and Marie-Odile Cordier (2001). “Supply Restoration in Power Distribution Systems — A Benchmark for Planning Under Uncertainty”. In: *Proceedings of the Sixth European Conference on Planning (ECP 2001)*. Ed. by Amedeo Cesta and Daniel Borrajo. AAAI Press, pp. 196–202.
- Thiébaux, Sylvie and Marie-Odile Cordier (2013). “Planning with MIP for Supply Restoration in Power Distribution Systems”. In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*. Ed. by Francesca Rossi. AAAI Press, pp. 2900–2907.
- Thiébaux, Sylvie, Jörg Hoffmann, and Bernhard Nebel (2005). “In Defense of PDDL Axioms”. In: *Artificial Intelligence* 168.1–2, pp. 38–69.
- Toda, Seinosuke (1991). “PP is as Hard as the Polynomial-Time Hierarchy”. In: *SIAM Journal on Computing* 20.5, pp. 865–877.
- Torralba, Álvaro (2015). “Symbolic Search and Abstraction Heuristics for Cost-Optimal Planning”. PhD thesis. Universidad Carlos III de Madrid.
- Torralba, Álvaro, Vidal Alcázar, Peter Kissmann, and Stefan Edelkamp (2017). “Efficient Symbolic Search for Cost-Optimal Planning”. In: *Artificial Intelligence* 242, pp. 52–79.
- Torralba, Álvaro, Stefan Edelkamp, and Peter Kissmann (2013). “Transition Trees for Cost-Optimal Symbolic Planning”. In: *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*. Ed. by Daniel Borrajo, Subbarao Kambhampati, Angelo Oddi, and Simone Fratini. AAAI Press, pp. 206–214.
- Torralba, Álvaro and Daniel Gnad (2017). *Alternatives to Explicit State Space Search: Symbolic Search*. Tutorial at the Twenty-Seventh International Conference on Automated Planning and Scheduling (ICAPS 2017).
- Torralba, Álvaro, Carlos Linares López, and Daniel Borrajo (2013). “Symbolic Merge-and-Shrink for Cost-Optimal Planning”. In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*. Ed. by Francesca Rossi. AAAI Press, pp. 2394–2400.
- Torralba, Álvaro, Carlos Linares López, and Daniel Borrajo (2016). “Abstraction Heuristics for Symbolic Bidirectional Search”. In: *Proceedings of the 25th*

- International Joint Conference on Artificial Intelligence (IJCAI 2016)*. Ed. by Subbarao Kambhampati. AAAI Press, pp. 3272–3278.
- van den Briel, Menkes, Romeo Sanchez, Minh B. Do, and Subbarao Kambhampati (2004). “Effective Approaches for Partial Satisfaction (Over-Subscription) Planning”. In: *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI 2004)*. AAAI Press, pp. 562–569.
- von Tschammer, Julian, Robert Mattmüller, and David Speck (2022). “Loopless Top-k Planning”. In: *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS 2022)*. Ed. by Sylvie Thiébaux and William Yeoh. To appear. AAAI Press.
- Yang, Bwolen, Randal E. Bryant, David R. O’Hallaron, Armin Biere, Olivier Coudert, Geert Janssen, Rajeev K. Ranjan, and Fabio Somenzi (1998). “A Performance Study of BDD-Based Model Checking”. In: *Proceedings of the Second International Conference on Formal Methods in Computer-Aided Design (FMCAD 1998)*. Ed. by Ganesh Gopalakrishnan and Phillip J. Windley. Springer, pp. 255–289.