

Versatile Cost Partitioning with Exact Sensitivity Analysis

Paul Höft¹, David Speck^{1,2}, Florian Pommerening², Jendrik Seipp¹

¹Linköping University, Sweden

²University of Basel, Switzerland

paul.hoft@liu.se, david.speck@liu.se, florian.pommerening@unibas.ch, jendrik.seipp@liu.se

Abstract

Saturated post-hoc optimization is a powerful method for computing admissible heuristics for optimal classical planning. The approach solves a linear program (LP) for each state encountered during the search, which is computationally demanding. In this paper, we theoretically and empirically analyze to which extent we can reuse an LP solution of one state for another. We introduce a novel sensitivity analysis that can *exactly* characterize the set of states for which a unique LP solution is optimal. Furthermore, we identify two properties of the underlying LPs that affect reusability. Finally, we introduce an algorithm that optimizes LP solutions to generalize well to other states. Our new algorithms significantly reduce the number of necessary LP computations.

Introduction

The objective of optimal classical planning is to find a cheapest sequence of actions that achieves some goal. One of the main methods for optimal planning is A* search (Hart, Nilsson, and Raphael 1968) with an admissible heuristic (Pearl 1984). Currently, the strongest admissible heuristics are based on *cost partitioning* (Katz and Domshlak 2010). In practice, cost partitions are often optimized using linear programs (LPs).

Saturated post-hoc optimization (Pommerening, Röger, and Helmert 2013; Seipp, Keller, and Helmert 2021) is a prominent example for this. It partitions costs according to a weighted sum of *saturated* cost functions. In each state encountered during the search these weights are optimized by an LP which incurs significant computational costs. Höft, Speck, and Seipp (2023) introduce methods to drastically reduce the number of solved LPs without compromising the quality of the heuristic. Their approach uses the concept of *sensitivity analysis* from Operations Research, which determines how changes to an LP affect its optimal solutions. In the context of (saturated) post-hoc optimization, the LPs computing cost partitions for two different states s and s' only differ in the bounds for some constraints. With sensitivity analysis one can often avoid computing a cost partition for s' by cheaply adapting the cost partition for s to s' .

In the work by Höft, Speck, and Seipp (2023), sensitivity analysis approximates the set of states where a cost parti-

tion can be adapted. Here, we go beyond that and develop a method that can characterize this set exactly.

Furthermore, we show that two properties of the underlying LP influence the relationship between an LP solution and the desired planning heuristic: *degeneracy*, where multiple solutions describe the same heuristic and *non-uniqueness*, where multiple heuristics are optimal. Improving the row formulation of the LP can reduce degeneracy, while improving the columns of the LP can reduce non-uniqueness.

In the presence of multiple equally good solutions, we would like to break ties in favor of solutions that generalize to more states. Therefore, we introduce an algorithm that favors LP solutions that correspond to more versatile cost partitions, increasing reusability. Our empirical evaluation shows that our new methods help to drastically reduce the number of LP computations required for the saturated post-hoc optimization heuristic.

Background

We consider SAS⁺ planning tasks (Bäckström and Nebel 1995) with action costs. The details of planning tasks are not important for this paper, the only relevant part is that a planning task induces a weighted transition system. A weighted *transition system* $\mathcal{T} = \langle S, L, T, cost, s_0, S_* \rangle$ consists of a finite set of *states* S , a finite set of *labels* L , a finite set of *labeled transitions* $T: s \xrightarrow{\ell} s'$ with $s, s' \in S$ and $\ell \in L$, a *cost function* $cost: L \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ that assigns a cost to each label, an *initial state* $s_0 \in S$, and a set of *goal states* $S_* \subseteq S$. A sequence of transitions leading from a state s to a goal state is called an *s-plan* and an s_0 -plan is called a *plan*. A plan is *optimal* if it has minimal cumulative cost.

Heuristic search (Bonet and Geffner 2001) is a common way of solving planning tasks optimally. A *heuristic function* $h: S \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ estimates the cost of the cheapest s -plan for each state s . It is *admissible* if it never overestimates the true cost $h_{\mathcal{T}}^*(s)$ of a cheapest s -plan. Strong representatives of admissible heuristics are *abstraction heuristics* (Edelkamp 2001; Helmert, Haslum, and Hoffmann 2007; Katz and Domshlak 2010; Seipp and Helmert 2018) which simplify the transition system of a planning task with a surjective *abstraction function* $\alpha: S \rightarrow S^\alpha$, yielding an *abstract transition system* $\mathcal{T}^\alpha: \langle S^\alpha, L, T^\alpha, cost, \alpha(s_0), S_*^\alpha \rangle$, where $T^\alpha = \{\alpha(s) \xrightarrow{\ell} \alpha(s') \mid$

$s \xrightarrow{\ell} s' \in T$ and $S_*^\alpha = \{\alpha(s) \mid s \in S_*\}$.

The preferable way of combining several heuristics admissibly is *cost partitioning* (Katz and Domshlak 2010; Pommerening et al. 2015). A cost partition C for a transition system \mathcal{T} is a tuple of cost functions $C = \langle cost_1, \dots, cost_n \rangle$, such that $\sum_{i=1}^n cost_i(\ell) \leq cost(\ell)$ for all labels $\ell \in L$. Evaluating a tuple of admissible heuristics $H = \langle h_1, \dots, h_n \rangle$ for state s under C yields the admissible estimate $h^C(s) = \sum_{i=1}^n h_i(cost_i, s)$, where each heuristic h_i is evaluated under cost function $cost_i$.

A strong heuristic based on cost partitioning is *saturated post-hoc optimization* (SPhO) (Pommerening, Röger, and Helmert 2013; Seipp, Keller, and Helmert 2021). For a transition system \mathcal{T} and a tuple of abstraction heuristics H for \mathcal{T} , the heuristic value $h^{\text{SPhO}}(s)$ for a state s is defined as the *objective value* of the linear program

$$\begin{aligned} \text{SPhO-LP}(s): \text{Minimize } & \sum_{\ell \in L} \text{cost}(\ell) \cdot Y_\ell \text{ subject to} \\ & \sum_{\ell \in L} \text{mscf}_h(\ell) \cdot Y_\ell \geq h(s) \quad \text{for all } h \in H \\ & Y_\ell \geq 0 \quad \text{for all } \ell \in L. \end{aligned}$$

Here, *mscf* is the *minimum saturated cost function* defined as $\text{mscf}(\ell) = \sup_{a \xrightarrow{\ell} b \in T} (h_{\mathcal{T}}^*(a) \ominus h_{\mathcal{T}}^*(b))$, where the \ominus operator is defined as regular subtraction for finite values. For infinite values, we use $x \ominus y = -\infty$ iff $x = -\infty$ or $y = \infty$, and $x \ominus y = \infty$ iff $x = \infty \neq y$ or $x \neq -\infty = y$ (Seipp, Keller, and Helmert 2020).

Linear programs (Thie and Keough 2008), such as SPhO-LP, can be written in the canonical form $\max_{x \in \mathbb{R}^n} \{c^\top x \mid Ax \leq b, x \geq 0\}$. Here x is the vector of n *decision variables*, $c \in \mathbb{R}^n$ is a vector of *objective coefficients*, $A \in \mathbb{R}^{m \times n}$ is the *coefficient matrix*, and $b \in \mathbb{R}^m$ is a vector of *bounds*. To solve a linear program with the simplex algorithm (Bradley, Hax, and Magnanti 1977), m slack variables are introduced to convert inequality constraints into equalities. The simplex algorithm explores solutions by iterating over valid *bases*. A *basis* partitions variables and slack variables into m *basic* and n *non-basic* variables. A valid partitioning requires invertibility of the $m \times m$ *basis matrix* B formed by the m columns of A associated with the *basic* variables. We indicate the indices of the basic variables as \mathcal{B} and the indices of the non-basic variables as \mathcal{N} . Each basis defines a unique *LP solution* by setting the non-basic decision variables to zero $x_{\mathcal{N}} = 0$, and the basic decision variables to $x_{\mathcal{B}} = B^{-1}b$ with objective value $c_{\mathcal{B}}^\top x_{\mathcal{B}}$, where $^{-1}$ represents the matrix inversion and $^\top$ the matrix transposition operation. Essential for this paper are the *reduced cost vector* $x_{\mathcal{B}}^*$, where $*$ indicates the optimality of the basis and the *inverse basis matrix* B^{-1} .

Exact Sensitivity Analysis

Höft, Speck, and Seipp (2023) establish that performing sensitivity analysis (Gal 1986) on the SPhO-LP can determine whether the computed cost partition for a state s can be cheaply adapted to another state s' without compromising the quality of the heuristic. They introduce four approaches,




Figure 1: Comparison of LP solver calls for the SPhO heuristic with exact and 100%-rule-based sensitivity analysis over PDBs for systematic patterns of size 1 and 2.

however, all of these approaches are approximations, so they may fail to recognize that a previously computed cost partition can be reused for a newly encountered state. The strongest among them is called “100% rule” and is guaranteed to require at most as many LP computations as any of the other methods (Note that despite its name, the 100% rule also is an approximation.)

A natural question is to what extent the number of LP computations can be further reduced if cost partitions are reused whenever possible. To perform such an exact sensitivity analysis, we build on the established result that a basis remains optimal under changes in its constraint bounds Δb if and only if all components in the reduced cost vector $x_{\mathcal{B}}^*$ remain greater than or equal to zero, i.e., $x_{\mathcal{B}}^* + \Delta x_{\mathcal{B}} \geq 0$ (Bazaraa, Jarvis, and Sherali 2009, Chapter 6). The effects on the reduced cost vector from changing the constraint bounds b by Δb are captured by the columns of the inverse basis matrix B^{-1} , as $\Delta x_{\mathcal{B}} = B^{-1}\Delta b$ (Vanderbei 2013, Chapter 7). Therefore, it holds that $x_{\mathcal{B}} + \Delta x_{\mathcal{B}} = x_{\mathcal{B}} + B^{-1}\Delta b$. Thus, changing the bounds b of an optimally solved LP by Δb will preserve optimality of the current solution as long as $x_{\mathcal{B}}^* + B^{-1}\Delta b \geq 0$.

Based on this result, we define the following algorithm that maintains a set of reduced cost vectors and inverse basis matrices $\langle x, B^{-1} \rangle$: When we encounter a state s during the search, we check whether one of the stored entries satisfies $x_{\mathcal{B}} + B^{-1}\Delta b \geq 0$. If so, we can use it to efficiently compute the heuristic value $h^{\text{SPhO}}(s)$. Otherwise, we solve SPhO-LP(s) and add its reduced cost vector and inverse basis matrix to our collection. We refer to Höft, Speck, and Seipp (2023) for the algorithm details, as our sensitivity analysis can be used in the same setting. We call this approach the *exact sensitivity analysis* for the h^{SPhO} heuristic since it allows us to compute h^{SPhO} while reusing a previous basis for *exactly* those states s where this is possible.

Figure 1 compares the number of LP solver calls required to compute the h^{SPhO} heuristic by exact sensitivity analy-

sis and the 100% rule.¹ This comparison is based on pattern databases (PDBs; Edelkamp 2001) for systematic patterns up to size 2 (Pommerening, Röger, and Helmert 2013). Overall, we observe a significant increase in the reusability of cost partitions when using exact sensitivity analysis. It increases the average percentage of state evaluations that do not require re-optimizing the LP from 79% to 85%. However, this advantage comes at the cost of slower evaluation and higher memory consumption, leading to a decrease in coverage from 818 to 682 solved tasks (out of 1827 tasks from the Optimal tracks of IPCs 1998–2018). The source code, benchmarks, and data for this and all following experiments are available online (Höft et al. 2024).

Alternative LP Solutions

All sensitivity analyses evaluate the reusability of a basis, not the reusability of the solution defined by the basis. There is a difference between these cases, because multiple bases can define the same solution coefficients x_B . This means that there can be multiple LP solutions that yield the same cost partition. A necessary condition for this is *degeneracy*. An LP solution is degenerate if its defining basis has basic variables set to 0 (Bazaraa, Jarvis, and Sherali 2009).

Proposition 1. *If an LP solution is degenerate there can be alternative solutions with the same solution coefficients.*

The opposite case are alternative solutions that describe a *different* cost partition, which is captured by the notion of LP solution *uniqueness* (Bazaraa, Jarvis, and Sherali 2009).

Proposition 2. *If an LP solution is non-unique there can be alternative solutions with different solution coefficients.*

So far, the topic of alternative optimal solutions for LP-based heuristics has not been discussed in the planning literature, because it is irrelevant when optimizing the LP for every state. However, when reusing LP solutions between different states, alternative solutions are of high interest, as their reusability can vary strongly. With the definitions of degeneracy and uniqueness, we can establish a precise relationship between sensitivity analysis and cost partitioning.

Theorem 1. *Given an LP solution sol for $SPhO-LP(s)$ inducing cost partition C , sensitivity analysis can determine the exact set of states for which C is optimal iff sol is non-degenerate and unique.*

Proof. Since sol is non-degenerate and unique, there are no other solutions for $SPhO-LP(s)$. Therefore, there are also no other cost partitions besides C for s that yield $h^{SPhO}(s)$. As a consequence, sol is reusable for exactly those states s' for which C is optimal. \square

We analyze the extent of degeneracy and uniqueness on $SPhO-LP$ solutions in Figure 2. It shows that the LPs computed for h^{SPhO} over different sets of pattern database heuristics have many alternative solutions. The number of degenerate solutions is highest for systematic patterns of sizes 1

¹Höft, Speck, and Seipp (2023) group abstractions with the same minimum saturated cost functions in a single constraint while Figure 1 uses one constraint per abstraction. We analyze abstraction grouping below.




Figure 2: Average percentage of degenerate and non-unique solutions for the SPHO heuristic over different PDB sets.

and 2. Removing patterns of size 1 from the collection decreases the average number of degenerate solutions slightly. A possible explanation for this is that the size-1 patterns are subsets of the size-2 patterns, which can lead to redundant constraints. Considering only size-1 patterns has much fewer degenerate solutions than the other two variants and even gives rise to some non-degenerate unique solutions.

Degeneracy and Non-Uniqueness

While a non-degenerate and unique LP solution describes a single optimal basis, degenerate or non-unique LP solutions imply that there are multiple optimal bases. The existence of multiple optimal bases can affect the reusability of a computed LP solution, since sensitivity analysis is defined for a specific optimal basis. As an example, consider three states: s_0 , s_1 , and s_2 . Further, assume that the SPHO LP for state s_0 yields two optimal bases, x_1 and x_2 , and performing a sensitivity analysis on x_i allows us to efficiently compute the heuristic value for state s_i . LP solvers can provide only one of these optimal bases at a time. So regardless of which basis (x_1 or x_2) the solver returns, an additional LP computation becomes necessary to derive the heuristic values for both states s_1 and s_2 . However, if we can reformulate the LP in a way that does not affect the computed heuristic value and at the same time reduces the space of optimal bases, ideally resulting in a single optimal basis (non-degeneracy and uniqueness), it becomes more likely that we can avoid the need for such additional LP computations.

We are therefore interested in reformulating the SPHO-LP to reduce the number of optimal bases by reducing degeneracy or non-uniqueness. Although it is generally not easy to predict whether a given LP will encounter degenerate solutions, there are two known criteria. Duplicate columns can lead to degenerate solutions and duplicate rows can lead to non-unique solutions (Bertsimas and Tsitsiklis 1997).

For the SPHO heuristic, this means that abstractions with the same minimum saturated cost function can affect degeneracy, and labels that have the same minimum saturated cost under all abstractions can affect non-uniqueness. Both observations are not new. Pommerening, Röger, and Helmert (2013) proposed to group duplicate labels, while

Algorithm 1 Greedily increase heuristic weights.

```

1: procedure INCREASEWEIGHTS( $\mathcal{H}, rem, s$ )
2:   for  $h \in \mathcal{H}$  with  $h(s) = 0$ , in random order do
3:      $\Delta w = \min_{\ell \in L} \left\{ \frac{rem(\ell)}{mscf_h(\ell)} \mid mscf_h(\ell) > 0 \right\}$ 
4:      $w_h += \Delta w$ 
5:     for  $\ell \in L$  do
6:        $rem(\ell) -= mscf_h(\ell) \cdot \Delta w$ 

```

Höft, Speck, and Seipp (2023) introduced abstraction grouping. Our analysis of optimal bases gives a novel explanation of why these techniques can reduce the number of required LP computations. Table 1 also shows empirical evidence that abstraction and label grouping are beneficial.

Finding Versatile Cost Partitions

As discussed above, there can be multiple cost partitions C that yield the same heuristic value $h^C(s) = h^{\text{SPhO}}(s)$ for a given state s . Therefore, instead of accepting the arbitrary cost partition C that the LP solver finds for state s , we hypothesize that it is beneficial to optimize C to obtain a more *versatile* cost partition C' . Such an optimized cost partition needs to preserve the estimate for s , i.e., $h^{C'}(s) = h^C(s)$, but apart from this requirement we can change it to make it optimal for more other states s' than the unoptimized C .

To obtain more versatile cost partitions, we turn to the dual of the SPhO-LP (Seipp, Keller, and Helmert 2021):

$$\begin{aligned} \text{Dual SPhO-LP}(s): \text{ Maximize } & \sum_{h \in \mathcal{H}} h(s) \cdot w_h \text{ s.t.} \\ & \sum_{h \in \mathcal{H}} mscf_h(\ell) \cdot w_h \leq cost(\ell) \text{ for all } \ell \in L \\ & w_h \geq 0 \text{ for all } h \in \mathcal{H}. \end{aligned}$$

Intuitively, this LP maximizes a weight w_h for each abstraction heuristic $h \in \mathcal{H}$. In the resulting cost partition C , each heuristic $h \in \mathcal{H}$ is assigned the cost function $cost_h$, where $cost_h(\ell) = mscf_h(\ell) \cdot w_h$. The value $h^{\text{SPhO}}(s)$ can then be computed as $h^{\text{SPhO}}(s) = h^C(s) = \sum_{h \in \mathcal{H}} w_h \cdot h(s)$.

When inspecting the SPhO-LP in dual form, it becomes apparent that the LP solver only optimizes the weights w_h for heuristics with a non-zero estimate $h(s)$. All heuristics h with $h(s) = 0$ do not factor into the optimization and their weights can be set arbitrarily by the LP solver. Therefore, we can increase the versatility of a cost partition C by increasing the weights w_h for heuristics h with $h(s) = 0$ to obtain cost partition C' . This will preserve the estimate for s , but possibly increase the estimates for other states s' , making it more likely that C' is optimal for s' than C . Additionally, increasing the weights can make the sensitivity analysis more versatile. This is the case since the new basis can only require fewer changes to be adapted for a new state s' .

We define a greedy procedure for increasing weights in Algorithm 1. It expects the remaining costs for all labels $\ell \in L$, which we compute as $rem(\ell) = cost(\ell) - \sum_{h \in \mathcal{H}} mscf_h(\ell) \cdot w_h$. Then the procedure iterates over the heuristics h with $h(s) = 0$ in a random order and increases weight w_h if the remaining label costs allow for an increase.

	Base	W+	GL	GH	GHL	GHL-W+
Base	–	45	89	58	83	87
W+	81	–	89	44	76	63
GL	116	102	–	85	28	40
GH	233	215	236	–	60	63
GHL	261	258	199	76	–	13
GHL-W+	276	256	222	96	34	–
Coverage	682	790	787	832	832	832

Table 1: Top: Per-task comparison of the number of LP solver calls for exact sensitivity analysis with different extensions. Cell (r, c) holds the number of tasks for which the algorithm in row r needs fewer LP solver calls than the algorithm in column c . The extensions are INCREASEWEIGHTS (W+) and grouping labels (GL), heuristics (GH), or both (GHL). Bottom: Total number of solved tasks.

If the INCREASEWEIGHTS procedure found a higher weight for at least one heuristic, we want to feed this solution back into the LP solver, so that its sensitivity analysis can benefit from the change. We do so by solving the same SPhO-LP again, but warm-starting it with the optimized weights. Re-solving the LP is cheap as long as the LP solver detects that the provided solution is already optimal.

If negative costs occur in the minimum saturated cost functions and the provided solution is degenerate or non-unique, the simplex algorithm may fail to detect its optimality, reject it, and search its own solution again (Bazaraa, Jarvis, and Sherali 2009). In our experimental analysis of SPhO with INCREASEWEIGHTS, we therefore skip the procedure when we detect negative saturated costs. This happens for roughly 30% of our benchmark tasks.

The number of times the INCREASEWEIGHTS procedure finds a weight to increase varies a lot between planning domains and even within tasks of one domain. In the Petri-Net-Alignment domain, for example, weights are increased between 0% and 21% of the time after finding an LP solution. As Table 1 shows, using INCREASEWEIGHTS usually decreases the number of LP solver calls and raises the number of solved tasks from 682 to 790 for the base case.

Conclusions

This paper focuses on the reusability of LP solutions computed for saturated post-hoc optimization. We introduced an exact sensitivity analysis that improves over existing approximations and empirically showed that it drastically reduces the required LP solver calls compared to the previous state of the art. Furthermore, we showed the importance of considering multiple alternative LP solutions and their potential to generalize to other states. Based on this insight, we proposed a novel greedy approach to optimize cost partitions for increased versatility, resulting in an even greater reduction in the number of LP computations required.

An interesting avenue for future work is to explore methods that not only aim to increase the versatility of a given LP solution, but also directly optimize for its reusability.

Acknowledgements

This work was partially supported by TAILOR, a project funded by the EU Horizon 2020 research and innovation programme under grant agreement no. 952215, and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at the National Supercomputer Centre at Linköping University partially funded by the Swedish Research Council through grant agreement no. 2022-06725. David Speck and Florian Pommerening were funded by the Swiss National Science Foundation (SNSF) as part of the project “Unifying the Theory and Algorithms of Factored State-Space Search” (UTA).

References

- Bäckström, C.; and Nebel, B. 1995. Complexity Results for SAS⁺ Planning. *Computational Intelligence*, 11(4): 625–655.
- Bazaraa, M. S.; Jarvis, J. J.; and Sherali, H. D. 2009. *Linear Programming and Network Flows*. John Wiley & Sons, Inc., 4th edition.
- Bertsimas, D.; and Tsitsiklis, J. N. 1997. *Introduction to Linear Optimization*, volume 6 of *Optimization and Neural Computation*. Athena Scientific.
- Bonet, B.; and Geffner, H. 2001. Planning as Heuristic Search. *Artificial Intelligence*, 129(1): 5–33.
- Bradley, S. P.; Hax, A. C.; and Magnanti, T. L. 1977. *Applied Mathematical Programming*. Addison-Wesley.
- Edelkamp, S. 2001. Planning with Pattern Databases. In Cesta, A.; and Borrajo, D., eds., *Proceedings of the Sixth European Conference on Planning (ECP 2001)*, 84–90. AAAI Press.
- Gal, T. 1986. Shadow Prices and Sensitivity Analysis in Linear Programming under Degeneracy. *OR Spectrum*, 8: 59–71.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.
- Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible Abstraction Heuristics for Optimal Sequential Planning. In Boddy, M.; Fox, M.; and Thiébaux, S., eds., *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling (ICAPS 2007)*, 176–183. AAAI Press.
- Höft, P.; Speck, D.; Pommerening, F.; and Seipp, J. 2024. Code and data for the ICAPS 2024 paper “Versatile Cost Partitioning with Exact Sensitivity Analysis”. <https://doi.org/10.5281/zenodo.10657675>.
- Höft, P.; Speck, D.; and Seipp, J. 2023. Sensitivity Analysis for Saturated Post-hoc Optimization in Classical Planning. In Gal, K.; Nowé, A.; Nalepa, G. J.; Fairstein, R.; and Rădulescu, R., eds., *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI 2023)*, 1044–1051. IOS Press.
- Katz, M.; and Domshlak, C. 2010. Optimal admissible composition of abstraction heuristics. *Artificial Intelligence*, 174(12–13): 767–798.
- Pearl, J. 1984. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley.
- Pommerening, F.; Helmert, M.; Röger, G.; and Seipp, J. 2015. From Non-Negative to General Operator Cost Partitioning. In Bonet, B.; and Koenig, S., eds., *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*, 3335–3341. AAAI Press.
- Pommerening, F.; Röger, G.; and Helmert, M. 2013. Getting the Most Out of Pattern Databases for Classical Planning. In Rossi, F., ed., *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, 2357–2364. AAAI Press.
- Seipp, J.; and Helmert, M. 2018. Counterexample-Guided Cartesian Abstraction Refinement for Classical Planning. *Journal of Artificial Intelligence Research*, 62: 535–577.
- Seipp, J.; Keller, T.; and Helmert, M. 2020. Saturated Cost Partitioning for Optimal Classical Planning. *Journal of Artificial Intelligence Research*, 67: 129–167.
- Seipp, J.; Keller, T.; and Helmert, M. 2021. Saturated Post-hoc Optimization for Classical Planning. In Leyton-Brown, K.; and Mausam, eds., *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*, 11947–11953. AAAI Press.
- Thie, P. R.; and Keough, G. E. 2008. *An introduction to linear programming and game theory*. John Wiley & Sons, 3rd edition.
- Vanderbei, R. J. 2013. *Linear Programming*. Springer-Verlag, 4th edition.