Research Statement: Heuristic Search for Domain-Independent Planning

Malte Helmert
Albert-Ludwigs-Universität Freiburg
Institut für Informatik
Georges-Köhler-Allee 52
79110 Freiburg, Germany
helmert@informatik.uni-freiburg.de

My main research interest is in heuristic search for domain-independent planning, in the setting where all actions are deterministic and the world is fully observable. This setting can be roughly described as a variant of "classical" heuristic search where the application domain is not known a priori to the algorithm designer, and hence issues which are tackled in application-dependent ways in classical search, most importantly the development of an appropriate heuristic function, have to be performed fully automatically based only on a declarative description of the transition semantics of the given problem (usually given in a logic-based language like PDDL).

I am interested in many aspects of heuristic search planning, ranging from purely theoretical studies to the development of practical planning systems. I am interested in satisficing (i.e., suboptimal) planning as well as optimal planning. In the following, I describe some topics in heuristic search planning that I have worked on recently.

- **novel heuristics for satisficing planning:**
  I have authored or co-authored the papers that introduced the *causal graph heuristic* (Helmert 2004), the *landmark heuristic* (Richter, Helmert, and Westphal 2008) and most recently the *context-enhanced additive heuristic* (Helmert and Geffner 2008), a generalization of the causal graph heuristic.
- **novel heuristics for optimal planning:**
  I have worked on novel approaches for finding good *pattern database heuristics* for planning (Haslum et al. 2007) and introduced *merge-and-shrink abstractions* (Helmert, Haslum, and Hoffmann 2007; 2008), also called *explicit-state abstractions*, which can be seen as generalizations of pattern database heuristics. Most recently, I have worked on a novel heuristic for optimal planning based on concepts drawn from landmarks and the classical max heuristic (not yet published).
- **planning systems:**
  I have developed the *Fast Downward* planning system (Helmert 2006a), in which most of the previously mentioned heuristics have been implemented. Fast Downward has won the classical track of the 4th International Planning Competition in 2004, and the LAMA planner built on top of Fast Downward has won the sequential satisficing track of the 6th International Planning Competition in 2008. I believe that researchers should strive towards openness in all of their work, and not just freely exchange ideas, but also engineering artifacts. The Fast Downward planner implementation has been used and extended in more than a dozen research projects both within our research group and at other research groups all over the world.
- **understanding planning search spaces:**
  I have worked on a number of papers that try to further our understanding of the characteristics of search spaces that make planning easy or hard. In particular, I have performed *domain-dependent complexity analyses* for the planning benchmark domains used in the 1998–2004 planning competitions (Helmert 2001; 2003; 2006b; Helmert, Mattmüller, and Röger 2006; Helmert 2008), providing an almost complete map of decision and approximation complexity for these benchmarks. More recently, inspired by the spectacularly bad performance of recent admissible planning heuristics in the Gripper benchmark domain, I have looked at the question how well an A^*-based planner could possibly perform when equipped with *almost perfect heuristics*, showing that in many classical planning domains optimal search with A^* can gain only very limited benefits over blind search even under very generous assumptions (Helmert and Röger 2008).
- **understanding planning heuristics:**
  Most recently, I have been very much interested in the question of *expressivity* of planning heuristics: what are the *fundamental strengths and limitations* of certain approaches to heuristic planning (such as relaxation, abstraction or use of landmarks)? How hard is it to come up with an *optimal* representative of a certain class of heuristics? Which *connections* exist between different classes of heuristics?

  My work on merge-and-shrink abstractions (Helmert, Haslum, and Hoffmann 2007) contains some discussion of such theoretical limitations for pattern databases, and this is also a central topic of a recent paper which compares the asymptotic accuracy of different planning heuristics in certain benchmark domains (Helmert and Mattmüller 2008).

  The paper on the context-enhanced additive heuristic

(Helmert and Geffner 2008) shows that there is a close relationship between the causal graph heuristic and additive heuristic, which were not previously considered to be very similar ideas.

Most recently, in as of yet unpublished work, I have been interested in establishing or disproving formal relationships between different classes of admissible planning heuristics (for example, landmark heuristics and abstraction heuristics).

References

Haslum, P.; Botea, A.; Helmert, M.; Bonet, B.; and Koenig, S. 2007. Domain-independent construction of pattern database heuristics for cost-optimal planning. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI-2007)*, 1007–1012. AAAI Press.

Helmert, M., and Geffner, H. 2008. Unifying the causal graph and additive heuristics. In *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS 2008)*, 140–147.

Helmert, M., and Mattmüller, R. 2008. Accuracy of admissible heuristic functions in selected planning domains. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*, 938–943. AAAI Press.

Helmert, M., and Röger, G. 2008. How good is almost perfect? In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*, 944–949. AAAI Press.

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.

Helmert, M.; Haslum, P.; and Hoffmann, J. 2008. Explicit-state abstraction: A new method for generating heuristic functions. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*, 1547–1550. AAAI Press.

Helmert, M.; Mattmüller, R.; and Röger, G. 2006. Approximation properties of planning benchmarks. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006)*, 585–589.

Helmert, M. 2001. On the complexity of planning in transportation domains. In Cesta, A., and Borrajo, D., eds., *Preproceedings of the Sixth European Conference on Planning (ECP'01)*, 349–360.

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

Helmert, M. 2004. A planning heuristic based on causal graph analysis. In Zilberstein, S.; Koehler, J.; and Koenig, S., eds., *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, 161–170. AAAI Press.

Helmert, M. 2006a. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Helmert, M. 2006b. New complexity results for classical planning benchmarks. In Long, D.; Smith, S. F.; Borrajo, D.; and McCluskey, L., eds., *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling (ICAPS 2006)*, 52–61. AAAI Press.

Helmert, M. 2008. *Understanding Planning Tasks – Domain Complexity and Heuristic Decomposition*, volume 4929 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag.

Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks revisited. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*, 975–982. AAAI Press.