

# Preface

Heuristic search is among the most widely used techniques in AI. In its different varieties, tree-based search and local search, it provides the core engine for applications as diverse as planning, parsing, and protein folding. One of the most promising avenues for developing improved search techniques is to integrate learning components that can adaptively guide the search. Research in this field is of wide interest in the AI community, not only because of the variety of subcommunities directly involved (problem solving, learning, constraint programming, operations research) but also because of the wide range of applications areas in which search algorithms play an important role.

Many disparate techniques have arisen in recent years that exploit learning to improve search and problem-solving. These techniques can be off-line or on-line, based on hard constraints or probabilistic biases, and applied to tree-structured or local search. This workshop brings together researchers and practitioners from the various subcommunities where such methods have arisen in order to learn from each other, develop common understandings, and inspire new algorithms and approaches.

Note that we are not discussing how search can be used to improve a learning algorithm, but rather the other way around. The 21 accepted papers discuss how the learning component can be employed to increase problem-solving performance, not just aid in scientific understanding. Relevant topics include, but are not limited to:

- adaptive and self-tuning algorithms
- automated parameter tuning

- automated portfolio design
- clause learning
- computing search space features
- decision-theoretic approaches to learning in search
- dynamic portfolio design
- exploiting models of search spaces
- exploiting performance profiles or run-time distributions
- incremental and active learning in search
- learning to select operators or heuristic functions
- model-based search
- reinforcement learning for search algorithms
- runtime prediction
- uncertainty in runtime prediction

This topic area is of particular interest at this time for two reasons. First, search itself is increasingly important. Most modern AI planning systems since the late 90s rely heavily on heuristic search, for instance, and search is also crucial in bioinformatics. The recent publication of Hoos and Stuetzle's textbook and the continued success of the SAT conferences are further evidence. Second, machine learning has matured to the point where efficient methods are widely available for handling the large amounts of data generated by search processes. Knowledge of how to apply these methods is also more widespread. Many isolated attempts to apply learning to improving search performance have been tried, dating back to the original founding of AI (GPS and SOAR, for example). However, to our knowledge there has been no recent venue where those interested in the topic could gather together to share ideas and experience.

Recent successful meetings that have addressed search include workshops at IJCAI, AAAI, and ICAPs. While many of the meetings have included elements that relate to learning for search from a particular perspective, none of them has focused exclusively on the topic, bringing together a wide range of perspectives at once.

We were pleasantly surprised by the number of submissions received. Each one was carefully reviewed by at least two members of the program committee. Due to scheduling constraints, only nine papers were selected for oral presentation, but all papers are represented in this technical report. We have contributions from the constraints community, from planning, on real-time search, and on learning heuristics, on local search, and on search trees. These papers truly represent a broad cross-section of the state of the art on learning for search.

*Wheeler Ruml and Frank Hutter*