

On-line search is driven by the need to commit to “actions” before their complete consequences are known. An “action,” in this context, can correspond to such diverse things as making a move in a two-player game, moving a robot, or allocating some resource (such as a page in a cache). On-line search can be necessary for a variety of reasons: there may be missing domain knowledge that has to be acquired actively, the domain may be known but so large that it cannot be searched completely in a reasonable amount of time, or it may simply be that the consequences of one’s actions depend on the behavior of some other entity. On-line search can also reduce the sum of planning and execution time.

The on-line search paradigm underlies many applications and has been independently investigated in

- artificial intelligence (for example, single-agent search and two-player games),
- robotics (for example, path planning and execution), and
- theoretical computer science,

among others. This has resulted in the development of a variety of on-line search approaches including assumptive planning, deliberation scheduling (including anytime algorithms), on-line algorithms and competitive analysis, real-time heuristic search, reinforcement learning, robot exploration techniques, and sensor-based planning.

The papers in these workshop notes were contributed by researchers from different fields who investigate on-line search approaches. They give an overview of the different methods, assumptions, and results, and hopefully enable the transfer of ideas between the different fields.

Sven Koenig (skoenig@cs.cmu.edu)  
Avrim Blum (avrim@cs.cmu.edu)  
Richard Korf (korf@cs.ucla.edu)  
Toru Ishida (ishida@kuis.kyoto-u.ac.jp)