

Reasoning About What to Plan

Richard Goodwin

School of Computer Science, Carnegie Mellon University
5000 Forbes Ave. Pittsburgh, Pennsylvania 15213-3890
rich@cs.cmu.edu

Agents plan in order to improve their performance. However, planning takes time and consumes resources that may in fact degrade an agent's performance. Ideally, an agent should only plan when the expected improvement outweighs the expected cost and no resources should be expended on making this decision. To do this, an agent would have to be omniscient. The problem of how to approximate this ideal, without consuming too many resources in the process, is the meta-level control problem for a resource bounded rational agent.

There are two central questions that have to be addressed for meta-level control: Where to focus planning effort and when to start executing the current best plan. These questions are interrelated. To start execution, the beginning of the plan must be elaborated to a level where it is operational. Even then, execution should only begin when the expected improvement due to further planning is outweighed by the cost of delaying execution. Once the agent has committed to executing some action, the planner can then disregard any plans inconsistent with this action and can concentrate on elaborating and optimizing the rest of the plan.

In my thesis research, I am exploring the use of sensitivity analysis based meta-level control for focusing computational effort. The object level problem of deciding which actions to perform is modeled as a standard decision problem and an approximate sensitivity analysis is performed. To facilitate the sensitivity analysis, actions, both abstract and operational, are augmented with methods for estimating their resource and time requirements. Methods are also needed to estimate the likelihood of events and action outcomes. All estimates include both the expected value and the expected range or variance. Information about the precision of estimates is critical when deciding whether to commit to a particular plan or whether to refine estimates through further computation or sensing.

When presented with a new task, the planner generates abstract plans for accomplishing the new and existing tasks. A sensitivity analysis identifies which of these plans are potentially optimal and non-dominated. Dominated and never-optimal plans are discarded. The sensitivity analysis also identifies which estimates the choice between plans is most sensitive to. Estimates that affect all plans more or less equally need not be refined. For instance, the occurrence of an earthquake may adversely affect all plans equally.

Determining the probability of an earthquake more exactly would not help in selecting between plans. Other factors may have differing effects. For instance, the likelihood of rain would help to choose between a plan to walk and a plan to drive somewhere. The sensitivity of a plan to particular estimates can also suggest ways of making the plan more robust. For instance, carrying an umbrella helps to reduce sensitivity to the likelihood of rain for the plan to walk.

When there are a number of plans that are potentially optimal and non-dominated and when the potential opportunity cost of selecting the wrong plan is significant, the meta-level controller directs the efforts of the planner to refine critical estimates. Estimates of resource use and action times can be improved by elaborating abstract operators into more operational operators or by simulated execution. Other object level estimates can be refined by adding more sensing to the plan or by additional computation using techniques such as temporal projection (Hanks 1990). Estimating computation time for complex planners is problematic. Further research is needed to determine how to best estimate and characterize expected plan improvement as a function of computation time.

Information from the sensitivity analysis and estimates of the cost of improving the current plan are used to make the tradeoff between the cost of delaying execution and the expected improvement in the plan for doing additional planning. Often systems that make this tradeoff ignore the fact that execution and planning can be overlapped in many situations. The DTA* algorithm is one example (Russell and Wefald 1991). In related work, I show how taking into account overlapping of planning and execution can improve performance (Goodwin 1994).

References

- Richard Goodwin. Reasoning about when to start acting. In K. Hammond, editor, *Proceedings of the Second International Conference (AIPS94)*. Artificial Intelligence Planning Systems, June 1994.
- Steve Hanks. Practical temporal projection. In *Proceedings, Eighth National Conference on Artificial Intelligence*. AAAI, July 1990.
- Stuart Russell and Eric Wefald. *Do the Right Thing*. MIT Press, 1991.