Preface

Metareasoning is the process of reasoning about reasoning itself. It is composed of both the meta-level control of computational activities and the introspective monitoring of reasoning. Meta-level control is the ability of an agent to efficiently trade off its resources between object level actions (computations) and ground level actions to maximize the quality of its decisions. While meta-level control allows agents to dynamically adapt their object level computation, it could interfere with ground level performance. Identifying the decision points that require meta-level control is of importance to the performance of agents operating in resource-bounded environments. Introspective monitoring is necessary to gather sufficient information with which to make effective meta-level control decisions. Monitoring may involve the gathering of computational performance data so as to build a profile of various decision algorithms. When reasoning fails at some task, it may involve the explanation of the causal contributions of failure and the diagnosis of the reasoning process.

Significant research questions also exist concerning the extent to which meta-level control and monitoring affects multiagent activity. We will compile these issues under the heading of multiagent metareasoning. In multiagent systems, where the quality of joint decisions affects individual outcomes, the value obtained by an agent exploring some portion of its decision space can be dependent upon the degree to which other agents are exploring complementary parts of their spaces. The problem of coordinated meta-level control refers to this question of how agents should coordinate their strategies to maximize the value of their joint actions.

Finally systems that address aspects of all three issues we term loosely models of self. A truly intelligent agent will have some conception of self that controls its reasoning choices, represents the products of monitoring, and coordinates the self in social contexts. Cox and Raja wrote a very brief manifesto that describes the above model of metareasoning and invited authors to compare and contrast their theories and implementations to this model. We also invited Don Perlis and Aaron Sloman to write invited papers on metareasoning and to present the results at the workshop.

The resulting papers contained in this collection fall into the four broad categories discussed above: meta-level control; introspective monitoring; multiagent metareasoning; and models of self. Each section contains both full papers and position statements. The significance of this workshop is that it brings together international researchers from different areas such as multiagent systems, planning/scheduling, logic, case-based reasoning and cognitive science to construct specific insights into technical and foundational issues related to these four topics. This goal fosters broad discussions of ongoing research, establishes directions for future research and joint collaborations, and identifies best practices for the evaluation of metareasoning systems.
.

– Michael T. Cox and Anita Raja, 13 May 2008