

Metareasoning: A Manifesto

Michael T. Cox and Anita Raja

Intelligent Computing, BBN Technologies, mcox@bbn.com
Department of Software & Information Systems, University of North Carolina at Charlotte, anraja@uncc.edu

Abstract

This manifesto proposes a simple model of metareasoning that constitutes a general framework to organize research on this topic. The claim is that metareasoning, like the action-perception cycle of reasoning, is composed of the introspective monitoring of reasoning and the subsequent meta-level control of reasoning. This model holds for single agent and multiagent systems and is broad enough to include models of self. We offer the model as a short conversation piece to which the community can compare and contrast individual theories.

Introduction

The 21st century is experiencing a renewed interest in an old idea within artificial intelligence that goes to the heart of what it means to be both human and intelligent. This idea is that much can be gained by thinking about one's own thinking. Traditionally within cognitive science and artificial intelligence, thinking or *reasoning* has been cast as a decision cycle within an action-perception loop similar to that shown in Figure 1. An intelligent agent perceives some stimuli from the environment and behaves rationally to achieve its goals by selecting some action from its set of competencies. The result of these actions at the ground level is subsequently perceived at the object level and the cycle continues. *Metareasoning* is the process of reasoning about this reasoning cycle. It consists of both the meta-level control of computational activities and the introspective monitoring of reasoning (see Figure 2). This cyclical arrangement represents a higher-level reflection of the standard action-perception cycle, and as such, it represents the perception of reasoning and its control.

The goal of *meta-level control* is to improve the quality of its decisions by spending some effort to decide what and how much reasoning to do as opposed to what actions to do. It balances resources between object level actions




Figure 1. The action-perception cycle

(computations) and ground level actions (behaviors). But while meta-level control allows agents to dynamically adapt their object level computation, it could interfere with ground level performance. Thus 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. It could involve generating explanations for object-level choices and their effect on ground level performance. When reasoning fails at some task, it may involve the explanation of the causal contributions of failure and the diagnosis of the object-level reasoning process.




Figure 2. Duality in reasoning and acting

Under the banner of *distributed metareasoning*, significant research questions also exist concerning the extent to which meta-level control and monitoring affects multi-agent activity. In multi-agent 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 any complete cognitive system that reasons about itself and its actions in the world will necessarily combine many aspects of metareasoning. 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. Hence a comprehensive approach will include *models of self* in support of metareasoning and integrated cognition.

Meta-Level Control

A significant research history exists with respect to metareasoning (Anderson and Oates 2007; Cox 2005), and much of it is driven by the problems of limited rationality. That is because of the size of the problem space, the limitations on resources, and the amount of uncertainty in the environment, finite agents can often obtain only approximate solutions. So for example with an anytime algorithm that incrementally refines plans, an agent must choose between executing the current plan or further deliberation with the hope of improving the plan. When making this choice, the agent is reasoning about its own reasoning (i.e., planning) as well as its potential actions in the world (i.e., the plan). As such this represents the problem of explicit control of reasoning.

Figure 2 illustrates the control side of reasoning along its upper portion. Reasoning controls action at the ground level in the environment; whereas metareasoning controls the reasoning at the object level. For an anytime controller, metareasoning decides when reasoning is sufficient and thus action can proceed. Although other themes exist within the metareasoning tradition (e.g., Leake 1996), this characterization is a common one (e.g., Raja and Lesser 2007; Hansen and Zilberstein 2001; Russell and Wefald 1991).

Now consider Figure 3. The most basic decision in classical metareasoning is whether an agent should act or continue to reason. For example the anytime planner always has a current best plan produced by the object level




Figure 3. Classical metareasoning (from Russell and Wefald 1991)

reasoning. Given that the passage of time itself has a cost, the metareasoner must decide whether the expected benefit gained by planning further outweighs the cost of doing nothing. If so it produces another plan; otherwise it executes the actions in the plan it already has. Note that this simple decision can be performed without reference to any perception of the ground level. Of course many more sophisticated meta-level control policies exist that include feedback.

Introspective Monitoring

The complementary side of metareasoning is less well studied. The introspective monitoring of reasoning about performance requires an agent to maintain some kind of internal feedback in addition to perception, so that it can perform effectively and can evaluate the results of metareasoning. For instance Zilberstein (Zilberstein and Russell 1996) maintains statistical profiles of past metareasoning choices and the associated performance and uses them to mediate the subsequent control and dynamic composition of reasoning processes.

But introspective monitoring can be even more explicit. If the reasoning that is performed at the object level (and not just its results) is represented in a declarative knowledge structure that captures the mental states and decision-making sequence, then these knowledge structures can themselves be passed to the meta-level for monitoring. For example the Meta-AQUA system (Cox and Ram 1999) keeps a trace of its story understanding decisions in structures called a Trace Meta-eXplanation Pattern (TMXP). Here the object-level story understanding task is to explain anomalous or unusual events in a ground-level story perceived by the system (see Figure 4).¹ Then if this explanation process fails, Meta-AQUA passes the TMXP and the current story representation to a learning subsystem. The learner performs an introspection of the trace to obtain an explanation of the explanation failure called an Introspective Meta-eXplanation Pattern (IMXP). The IMXPs are used to generate a set of learning goals that are passed back to control the object-level learning and hence improve subsequent understanding. TMXPs explain *how* reasoning occurs; IMXPs explain *why* reasoning fails.




Figure 4. Introspective monitoring in Meta-AQUA

Note that the object-level process described above is a story understanding task without reference to the execution

¹ Meta-AQUA does no action at the ground level. Rather it perceives events representing characters in the story doing actions.

of personal actions at the ground level. The emphasis here is upon the perception and monitoring side of the model; that is, the understanding or comprehension processes in the model are equally as important as the action and control processes were in Figure 3, and indeed they can be treated independently. However most systems, especially agent-based systems, combine both in various fashions.

Distributed Metareasoning

In a multi-agent context, if two or more agents need to coordinate their actions, the agents' meta-control components must be on the same page. The agents must reason about the same problem and may need to be at the same stage of the problem-solving process. For example, suppose one agent decides to devote little time to communication/negotiation (Alexander, Raja, Durfee, and Musliner 2007) before moving to other deliberative decisions while another agent sets aside a large portion of deliberation time for negotiation; the latter agent would waste time trying to negotiate with an unwilling partner.

We define an agent's problem solving context as the information required for deliberative-level decision making, including the agent's current goals, action choices, its past and current performance, resource usage, dependence on other agents, etc. Suppose the agent's context when it is in the midst of execution is called the current context, while a pending context is one where an agent deliberates about various what-if scenarios related to coordination with other agents. Distributed metareasoning can also be viewed as a coordination of problem solving contexts. One meta-level control issue would be to decide when to complete deliberation in a pending context and when to replace the current context with the pending context. Thus if an agent changes the problem solving context on which it is focused, it must notify other agents with which it may interact. This suggests that the meta-control component of each agent should have a multi-agent policy where the content and timing of deliberations are choreographed carefully and include branches to account for what could happen as deliberation (and execution) plays out. Figure 5 describes the interaction among the meta-level control components of multiple agents.

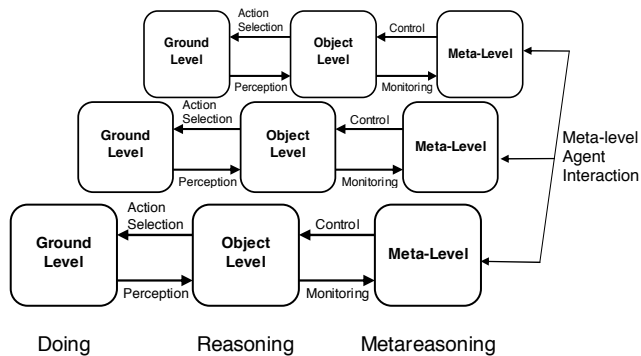


Figure 5. Meta-level reasoning among multiple agents

Another meta-control question when there are multiple pending contexts is to determine which pending context should be allocated resources for deliberation. In all of these examples, the metareasoning issues are a superset of single agent cases.

Models of Self

For a cognitive agent to behave intelligently in a physical and social environment with complex, dynamic interactions, many if not all of the features necessary for an integrated human-level model of intelligence are required. For it to succeed in such environment, an agent must perceive and interpret events in the world including actions of other agents, and it must perform complex actions and interact in a social context. These constitute the minimal object level requirements. At the meta-level, an agent must have a model of itself to represent the products of experience and to mediate the choices effectively at the object level. Facing novel situations the successful agent must learn from experience and create new strategies based upon its self-perceived strengths and weaknesses. Consider Figure 6.




Figure 6. Example integrated model of self

Monitoring at the meta-level can determine the kinds of mental actions at which the agent excels and those it fails. Using such introspective information allows the agent to choose reasoning strategies that best fit future intellectual demands like the agent that selects actions based on past task performance. In more complicated approaches, the agent may actually construct a complex reasoning strategy rather than simply choose an atomic one. In either case, the basis for such metareasoning comes from a picture of itself, its capacities (both physical and mental), and its relationships to other agents with which it must interact to recognize and solve problems.

Many theorists have speculated as to the interactions between levels of representation and process (i.e., the architecture), but few researchers have attempted to implement the full spectrum of computation implied in a comprehensive model of self (see Singh 2005, for one such attempt). However we challenge the AI community to consider seriously the problems of metareasoning in this larger context. How would an agent best understand itself and use such insight to construct a deliberate knowledge-level reasoning policy? Can an agent know enough about

itself and its colleagues' self-knowledge to communicate its meta-level needs for coordination? Can it estimate the time it might take to negotiate a coordination policy with its fellow agents and hence negotiate the time and length of a negotiation session? Finally could an intelligent soccer agent decide that it is good at planning but getting weak at passing and so aspire to becoming a coach? We claim that the model of acting, reasoning, and metareasoning put forth in this document can help maintain clarity if this challenge is to be embraced and answering questions like these pursued.

Conclusion

This manifesto has tried to present in plain language and simple diagrams a brief description of a model of metareasoning that mirrors the action-selection and perception cycle in first-order reasoning. Many theories and implementations are covered by this model including those concerning meta-level control, introspective monitoring, distributed metareasoning, and models of self. We claim that it is flexible enough to include all of these metacognitive activities, yet simple enough to be quite parsimonious. Figures 3 through 6 and their accompanying examples suggest some variations on the potential implementations rather than dictate an agenda. We offer the model as a framework to which the community can compare and contrast individual theories, but most of all, we hope that this model can clarify our thinking about thinking about thinking.

References

- Alexander, G.; Raja, A.; Durfee E.; and Musliner, D. 2007. Design Paradigms for Meta-Control in Multi-Agent Systems In A. Raja and M. T. Cox eds., *Proceedings of AAMAS 2007 Workshop on Metareasoning in Agent-Based Systems* 92-103. Hawaii, May.
- Anderson, M. L., and Oates, T. 2007. A Review of Recent Research in Metareasoning and Metalearning. *AI Magazine* 28(1): 7-16.
- Cox, M. T. 2005. Metacognition in Computation: A Selected Research Review. *Artificial Intelligence* 169(2): 104-141.
- Cox, M. T., and Ram, A. 1999. Introspective Multistrategy Learning: On the Construction of Learning Strategies. *Artificial Intelligence*, 112, 1-55.
- Hansen, E., and Zilberstein, S. 2001. Monitoring and Control of Anytime Algorithms: A Dynamic Programming Approach. *Artificial Intelligence* 126(1-2):139-157.
- Leake, D. B. 1996. Experience, Introspection, and Expertise: Learning to Refine the Case-Based Reasoning Process. *Journal of Experimental and Theoretical Artificial Intelligence* 8(3): 319-339.
- Raja, A., and Lesser, V. 2007. A Framework for Meta-Level Control in Multi-Agent Systems. *Autonomous Agents and Multi-Agent Systems* 15(2): 147-196.

Russell, S. J., and Wefald, E. 1991. Principles of Metareasoning. *Artificial Intelligence* 49: 361-395.

Singh, P. 2005. EM-ONE: An Architecture for Reflective Commonsense Thinking. Ph.D. diss., Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Boston, Mass.

Zilberstein, S., and Russell, S. J. 1996. Optimal Composition of Real-Time Systems. *Artificial Intelligence* 82(1-2):181-213.