Metacognition for Multi-Agent Systems *

Rajiv T. Maheswaran and Pedro Szekely
University of Southern California - Information Sciences Institute
4676 Admiralty Way #1001, Marina Del Rey, CA 90292
{maheswar,pszekely}@isi.edu

Abstract

We consider the problem of determining metacognition strategies for multi-agent systems. A single-agent working in isolation can choose its mode of operation solely based on environmental conditions. However, an agent in a team must consider the modes of other agents as the information available to it and thus its decision process depends critically on the information and decision process of others. We present a model and discuss an example that illustrates the complexity of the problem.

1 Introduction

While *metacognition* has many diverse definitions under the "thinking about thinking" umbrella [1], we consider the interpretation of it as the monitoring and regulation of cognitive processes. In computational agents that focus on decision-making, this manifests itself in choosing between multiple reasoning strategies and determining the amount of resources (computation, memory, bandwidth, etc.) to use in pursuing each strategy . These choices can be interpreted as an agent choosing one or more *modes* of operation.

This mode is often based on environmental conditions. Internal environmental conditions could capture properties such as the hardware capabilities of the agent. This is particularly important for mobile agents who may wish to modify their processes based on the computation, memory and communication available to them. External environmental conditions could include characteristics of the problem instantiation being faced, such as the volatility/stability of the domain or the role required to be performed by the agent.

*The work presented here is funded by the DARPA COORDINATORS Program under contract FA8750-05-C-0032. The U.S. Government is authorized to reproduce and distribute reports for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of any of the above organizations or any person connected with them.

A single-agent can monitor its environment and choose a decision process dynamically based on a metacognition strategy that was obtained *a priori*. One path to developing such a strategy is by learning optimal choices over multiple environments. This procedure is insufficient in a multi-agent team because system performance does not depend solely on environmental factors. Agents decisions are made based on the information available. In multi-agent systems, the information available and how often it is updated depends on the decision processes chosen by other agents. Thus, the mode of operation of one agent chosen by its metacognition strategy affects and is affected by the modes of operation chosen by the other agents in the team.

In multi-agent systems, performance is determined by a *joint* set of modes. The metacognition goal should be to find a way to choose this joint set of modes such that the resulting behavior is good. This choice may be the result of a decentralized strategy that yields a consensus on a common mode or heterogeneous distribution of modes such that any penalties for heterogeneity are offset by its gains.

2 Model

Here, we discuss a formal model of single and multi-agent systems and furthermore, what metacognition means within these contexts. Illustrations of the models and parameters can be seen in Figures 1 and 2. A single agent gets information from the environment (η_1^E), which is transformed (f_1^{EI}) into some internal information (η_1^I). This internal information serves as the base input into a decision process (f_1^{IA}) by which an action (a_1) is chosen. The agent has a choice to determine the transformation of environmental information, the internal information and decision process from a larger set of possibilities: $\eta_1^E \in H_1^E, f_1^{EI} \in F_1^{EI}, f_1^{IA} \in F_1^{IA}$. Metacognition for a single agent can be represented as choosing an element of $M_1 := H_1^E \times F_1^{EI} \times F_1^{IA}$ based on information from the environment. The resulting element is the mode of operation that an agent chooses.

In a multi-agent system, the internal information from an agent (η_i^I) is also a function of the information that it receives from other agents ($\eta_i^j, j \neq i$). This requires looking at the problem in a new way. If agents come up with metacognition strategies (methods of picking modes) individually, then they must consider the information from other agents ($\{\eta_i^j\}_{j \neq i}$) as additional environmental input. This would lead to uncoordinated metacognition strategies in the multi-agent system. A better way would be to choose an element from $M_1 \times M_2 \times H_2^1 \times H_1^2$ that optimizes the team behavior based on the joint environment (η_1^E, η_2^E). The challenge would be to achieve this joint strategy in a distributed manner. The motivation for developing such an approach is provided in an example that follows.

3 Example

Consider a situation where two agents are trying to coordinate their activities over an extended period of time [2]. They can be in one of two modes of operation. In the *fast-planning* mode, the agent has determined that either the environment is too volatile for

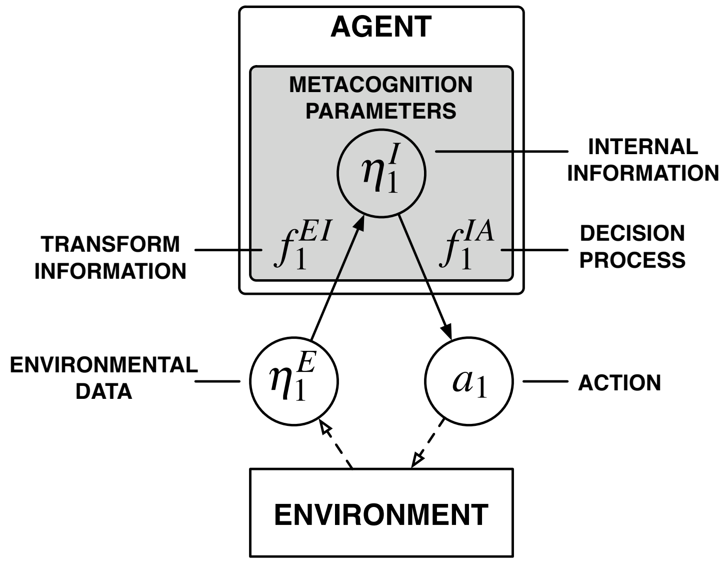

Figure 1: Metacognition Model and Parameters for a Single Agent

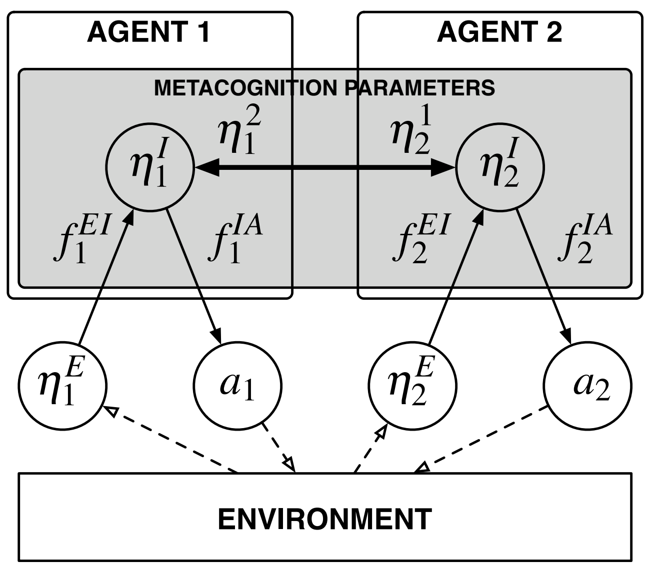

Figure 2: Metacognition Model and Parameters for a Multi-Agent System

extended analysis or that it does not have the capability to do complex deliberation. Thus, it sends small messages very often about its status and immediate plans. In the *long-planning* mode, an agent has determined that the environment and its capabilities can support deeper reasoning or it may only have periodic access to communication facilities. In this situation, the agent considers all the possible state evolutions and contingent actions over a longer time horizon. It sends larger messages with details of this information but not as often.

If the agents choose their metacognition strategy based solely on environmental conditions, they may end up in different modes. Then, most of the messages of the fast-planning agent will be ignored or provide insufficient information for the long-planning agent. The messages of the long-planning agent will be too cumbersome and infrequent for the fast-planning agent.

While the choices may have seem appropriate for the environment, they may lead to worse performance by wasting resources due to lack of coordination. If the long-planning agent had chosen a fast-planning strategy despite his environment, the team might have performed better because it would be providing useful information to the other agent. Also, if the fast-planning agent had incorporated the fact that the other agent was using a long-planning strategy, it might have been able to use some of the cycles it had spend to construct messages that were going to be ignored on more fruitful calculations. Thus, if agents feature a metacognition component, this component must engage in *coordinated metacognition* to be successful.

4 Summary

We discussed how metacognition reasoning differs between a single and multi-agent system. The primary complexity is that the information an agent receives, in a multi-agent system, depends on the information provided by other agents in the system. Thus, a model of metacognition must address the coupling that exists between the choices of modes of operations. We provide an example scenario that illustrates how coordinated metacognition strategies may be more useful that those that appear optimal in isolation. The challenge is to develop methods to achieve coordination at the meta-level in a distributed manner.

References

[1] Michael T. Cox. Metacognition in computation: a selected research review. *Artificial Intelligence Journal (AIJ)*, 169(2):104–141, 2005.

[2] Rajiv T. Maheswaran, Craig M. Rogers, Romeo Sanchez, and Pedro Szekely. Distributed coordination in uncertain multiagent systems. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS 2007)*, Honolulu, HI, May 2007.