Modelling of Agents' Behavior with Semi-collaborative Meta-agents

Jan Tožička, Filip Železný, and Michal Pěchouček

Czech Technical University,
Gerstner Laboratory, Prague, Czech Republic
{tozicka, zelezny, pechouc}@labe.felk.cvut.cz
http://gerstner.felk.cvut.cz

Abstract. An autonomous agent may largely benefit from its ability to reconstruct another agent's reasoning principles from records of past events and general knowledge about the world. In our approach, the meta-agent maintains a first-order logic theory, called the community model, yielding predictions about other agents' decisions. In this contribution we introduce a query-based collective reasoning process where the semi-collaborative meta-agents use active learning technique to improve their models. We provide empirical results that demonstrate the viability of the concept and show the benefits of collective meta-reasoning.

1 Introduction

We are working with a community of autonomous reasoning agents endowed by number of capabilities which allow them to form coalitions to solve complex tasks (e.g. logistics). Behavior of our agents is given by a permanent reasoning algorithm and a set of private knowledge describing e.g. agent's preferences concerning the tasks, cooperation with other agents, etc. Agent's private knowledge is permanent. Dynamics of agent's behavior is given by changing resources and their availability, by ever changing environment in which the agents operate, and by different behavior of the other members of the multi-agent community.

What can also change or evolve during the lifespan of an agent is agent's awareness about the private knowledge of other members of the community. An agent has no direct access to the private knowledge of any other agent, it only can try to estimate or reconstruct its content in order to e.g. influence complexity, quality and effectiveness of collaboration, as well as the response time of the system.

Meta-reasoning is a key concept in this article. Unlike in classical computer science literature [1], where the meta-reasoning process is strictly understood as a reasoning process about yet another reasoning process, we will refer to meta-reasoning as agent's capability to reason also about other agent's knowledge, preferences, etc. In this contribution we compare deductive vs. inductive approach to meta-reasoning and introduce collective meta-reasoning of semi-collaborative meta-agents.

In Section 2, we will firstly introduce used meta-reasoning architecture and then present possible approaches to collective meta-reasoning. Section 3 presents a logistic scenario and experiments evaluating presented methods. We summarize our contribution in Section 4.

M. Pěchouček, P. Petta, and L.Z. Varga (Eds.): CEEMAS 2005, LNAI 3690, pp. 572–575, 2005.

2 Meta-reasoning Architecture

In this section, we will briefly introduce meta-agents and collective meta-reasoning. More detailed description of architecture, knowledge, formal language and used logic can be found in [2]. The principal role of the meta-reasoning agent is to support meta-reasoning process through maintaining and exploiting a model of the agent community. Our meta-reasoning agent monitors the community and is completely independent from the functionality of the community of the 'ordinary' agents. In principle, the community model can be maintained in two ways: (i) **deductive reasoning** maintains the model to contain only knowledge that logically follow from the observations. It is implemented by resolution-principle-based automated prover; and (ii) **inductive reasoning** maintains an *approximative model* which also contains knowledge generalizing the monitored knowledge; this formula can prove to be in conflict with some future events. Inductive meta-reasoning has been implemented using inductive logic programming (ILP).

The meta-agent can use different AI methods in order to update assumed model to consider new incoming events, we will call this process **model revision** operation, and to query assumed model during **model inspection** operation, which has three possible outcomes for given query: yes, no, and $unsure$.

2.1 Collective Meta-reasoning

Our experimental configuration characterized by a set of agents observed by a set of deductive/inductive meta-agents naturally lends itself as a ground on which some interesting, recently emerged machine-learning approaches can be empirically evaluated. Specifically, we have taken inspiration from (i) the study [3] on distributed learning of first-order logic theories; (ii) the *active learning* framework, where the learner is allowed to actively pose queries to an oracle; and (iii) the paradigm of *closed-loop* learning [4], where the learner can initiate experiments determining the required answer to a query. Adapting these techniques in the agent environment, has the promise of achieving a favorable trade-off between the average quality of the models developed by the meta-agents, and the invested computational effort.

Indeed, pursuing the outlined efficiency motivation, the paper [3] demonstrates that the search for a first-order logic theory. Collaboration can take place in the query time: a query is answered by several agents and the collective answer may be obtained by voting – we call this approach **deductive collaboration**. However, we try to establish interaction in the inductive process itself. In active learning, the learner is able to actively pose queries to an oracle, whose answers guide the model formation. Our adaptation of this principle into an **inductive collaboration** scheme assumes that an inductive meta-agent, besides the ability to generalize provided learning examples, may query another meta-agent, whose answer follows from its current model (possibly only partially built). Query is created randomly even if several heuristic approaches have been identified.

In the collective meta-reasoning development we will also apply the ideas of *closed-loop* learning [4], where the learner actually triggers experiments determining the required answer to a query. In our scenario, a meta-agent, monitoring an agent A, collaborates with another agent B, from whom it asks to offer a coalition to A. The B's proposal along with the proposal outcome then form an new observations.

3 Experiments

We have experimented with our meta-reasoning ideas in the $\mathcal{A}$-**cross** multi-agent scenario that has been integrated in the $\mathcal{A}$-**globe** multi-agent platform [5]. For us, the most important part of $\mathcal{A}$-**cross** logistic scenario are **transport-agents** who organize the transport of commodities. They form coalitions in order to convey the cargo. Coalition formation of transport-agents is determined by (i) availability of resources and (ii) sets of collaboration restrictions. We have extended this scenario by **observer-agents** and **meta-agents** that implement our meta-reasoning architecture. The observer-agents watch the transport-agents in their neighborhoods, transform observations into formal language and send them to the meta-agents. The meta-agents build their models about the community and try to reconstruct collaboration restrictions of transport-agents.

The meta-reasoning in our scenario works with *events* and *queries*. Events are created based on messages sent during CNP communication protocol. Queries can be used by user to get some knowledge from created community model. Queries asked among the meta-agents (as described in the section 2.1), has similar form as events. The meta-agents are semi-collaborative as they can agree to cooperate with other meta-agents (depending on their private restrictions) and even if one decides to cooperate he will answer only limited number of other agents queries.

Our goal is to evaluate the quality of a meta-reasoning process in different *configurations* of the interactions between meta-agents. A configuration is described by a rooted directed graph, where vertices correspond to meta-agents and edges lead from a query-posing agent to the answering agent (an oracle). Each possible configuration is characterized by two parameters, called **distributedness** and **collaborativeness**. The distributedness of a configuration graph is defined as the average distance among all pairs of vertices in the graph. The collaborativeness is defined as the maximum number of agents querying the same oracle (i.e.. the maximum branching factor).

The *quality* of a meta-reasoning process is viewed as a trade-off between the total computational effort used in model developing and the average model quality achieved. We will proceed by fixing a total budget and measuring the average model quality. Here, the quality of a created model is its **predictive accuracy**, i.e.. the ration of correctly classified *test* observations containing observations that are not used for learning.

To postulate expense-consciousness among the meta-agents, we establish a form of information market in the community, with the following rules of trade: (i) a meta-agent has to pay for every yes or no answered query; (ii) adhering to a common understanding of *information value*, the price for a answered query should be low if so is the model quality used to answer the query; and (iii) the meta-reasoning process initiates by assigning a constant budget to each agent. Each agent adds to its budget any price it charges to another agent, and subtracts any price it is charged by another agent. The meta-reasoning process terminates when all agents' budgets have been consumed.

Results. Figure 1 shows the average model quality values in respect to the distributedness values for 35 randomly generated configurations. It is interesting to note that the trend-line fitting the average quality of the models grows with the value of distributedness. Similarly, Figure 2 shows that the average model quality grows with the collaborativeness, as could be expected, however the growth is remarkably slow.

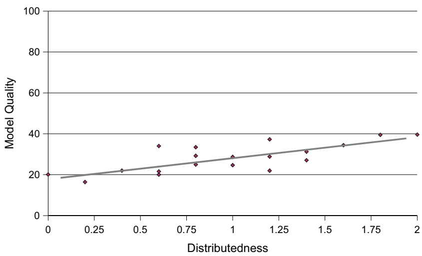

Fig. 1. X-Axis shows the distributedness of meta-agents and Y-Axis shows the average quality of the models

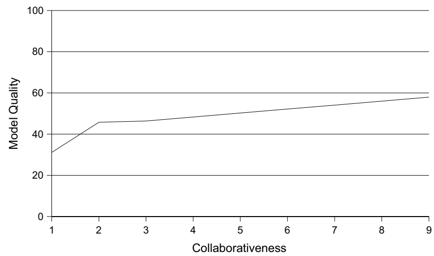

Fig. 2. X-Axis shows the collaborativeness among meta-agents and Y-Axis shows the average quality of their models

Both experiments suggest that distributed configurations, where only few meta-agents have direct or close access to the observations, allow a high average quality of the models created by the meta-agents, and, in the frame of a simple resource-conscious framework, they seem even superior to the centralized configurations.

4 Conclusions

In this paper, we focus on cooperation within a group of semi-collaborative meta-agents. Implemented technique in $\mathcal{A}$-**cross** scenario is evaluated in respect to the *distributedness* and the *collaborativeness* of meta-agents. The goal of meta-reasoning was to predict, based on the previous observations or using active learning technique, whether an agent will agree to join a coalition. Collective meta-reasoning proved to be useful when the meta-agents have to solve the trade-off between the average model quality and the total invested effort. We have shown, that under suitable settings, more distributed configuration can bring better model qualities than centralized case.

Acknowledgement. The presented research has been in parts supported by Office for Naval Research, project no.: N00014-03-1-0292 and European Office for Aerospace Research and Development, project no. F61775-99-WE099.

References

1. Maes, P.: Computational reflection. Tech. report 87-2, University of Brussels, AI Lab (1987)
2. Tožička, J., Bárta, J., Pěchouček, M.: Meta-reasoning for agents' private knowledge detection. In Klusch, M., Ossowski, S., Omicini, A., Laamanen, H., eds.: Cooperative Information Agent VII – Lecture Notes in Computer Science, LNAI 2782, Springer-Verlag, Heidelberg (2003)
3. Železný, F., Srinivasan, A., Page, D.: A monte carlo study of randomised restarted search in ilp. In: Inductive Logic Programming, 14th Int. Conf., Berlin, Springer (2004) 341–358
4. King, R., Whelan, K., Jones, F., Reiser, P.: Functional genomic hypothesis generation and experimentation by a robot scientist. Nature **427** (2004) 247–252
5. Šišlák, D., Rollo, M., Pěchouček, M.: A-globe: Agent platform with inaccessibility and mobility support. In Klusch, M., ed.: CIA VIII. Number 3191 in LNAI, Springer-Verlag (2004)