Towards a Cognitive Meta-Model for Adaptive Trust and Reputation in Open Multi-Agent Systems

(JAAMAS Extended Abstract)

Bruno W. P. Hoelz
University of Brasília
bhoelz@unb.br

Célia G. Ralha
University of Brasília
ghedini@unb.br

ABSTRACT

Computational trust and reputation models are key elements in the design of open multi-agent systems. They offer means of evaluating and reducing risks of cooperation in the presence of uncertainty. However, the models proposed in the literature do not consider the costs they introduce and how they are affected by environmental aspects. In this paper, a cognitive meta-model for adaptive trust and reputation in open multi-agent systems is presented. It acts as a complement to a non-adaptive model by allowing the agent to reason about it and react to changes in the environment. We demonstrate how the meta-model can be applied to existent models proposed in the literature, by adjusting the model's parameters. Finally, we propose evaluation criteria to drive meta-level reasoning considering the costs involved when employing trust and reputation models in dynamic environments.

Keywords

open multi-agent systems, meta-model, adaptiveness, computational trust and reputation

1. INTRODUCTION

In a multi-agent system (MAS), autonomous agents often need to cooperate to achieve their goals. Consequently, agents need to trust each other, explicitly or not. This decision can be based on previous, direct experience or by obtaining information about the reputation of potential partners. While many trust and reputation (T&R) models with different characteristics have been proposed in the literature, from the standpoint of an autonomous agent, there is no evaluation criteria to guide its decision on which model to use and how to adapt it to the dynamic environment of an open MAS.

2. THE AGENT MODEL

In order to allow the agent to reason about its T&R model and adapt it to changes in the environment, we present a cognitive meta-model for T&R adaptation. The agent model containing the meta-model, shown in Figure 1, is divided in two levels: the meta-level, consisting of abstract representations of environmental aspects and common elements of T&R; and the implementation level, that includes the concrete, model-specific realization of these elements.

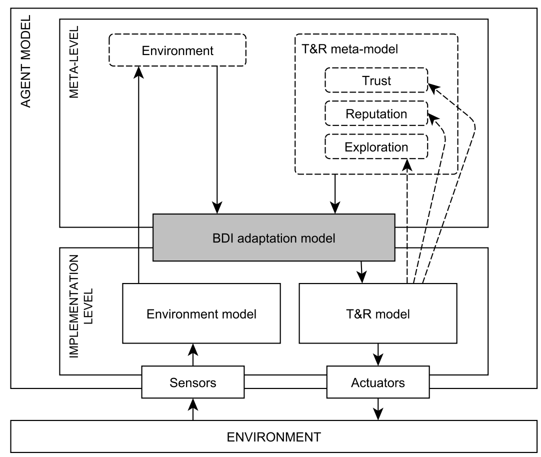

Figure 1: Agent model for T&R adaptation

The environment element of the meta-level is a meta-model that represents domain-independent aspects used in the adaptation process. These include operating costs, frequency of transactions, earned and total utility per unit of time, availability of trusted partners, availability of information sources, communication costs, and cost of information acquisition. Domain-specific aspects can be mapped to one or more meta-level aspects and later in the adaptation process they can be reintroduced via the agent's adaptation model.

The T&R meta-model is divided into three sub-models that are subject to adaptation: trust, reputation and exploration. The T&R sub-models define their respective information sources and how they are used in the T&R evaluations. The exploration sub-model defines how the agent initiates its interactions in the absence of previous information and how it seeks new partners and sources. Figure 2 presents the components of each sub-model.

One of the main characteristics of a T&R model is the set of information sources used in the evaluation of trust. The two most common sources are direct experience and witness information, while others such as direct observation, bias (or prejudice), and norms have been proposed in the literature [1]. In the T&R meta-model, an information source (IS)

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

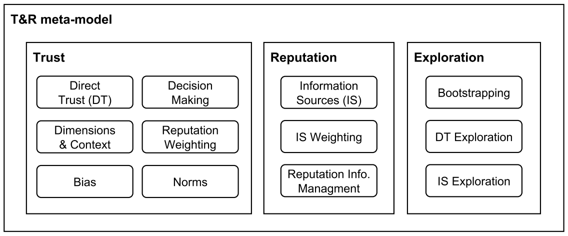

Figure 2: Components of the T&R meta-model

is composed of elements that determine how the agent (i) acquires and shares information, (ii) manages its memory, and (iii) evaluates the reliability and credibility of the IS. These elements have a direct relation with the environmental conditions. For example, information acquisition may be affected by communication costs, and memory management may have to be adjusted to the frequency of transactions.

3. BDI ADAPTATION MODEL

The adaptation model is a Belief-Desire-Intention (BDI) model that connects the meta-level to the implementation level. It allows the agent to reason about T&R adaptation. Beliefs are used to represent facts about the environment and the T&R model, as represented by their meta-level counterparts. To represent the agent's desires, two types of goals are used: monitoring and adaptation goals. Monitoring goals signal the need to adjust the model in use. They trigger an adaptation event that results in the selection of adaptation goals. Next, a plan selection process evaluates the set of available adaptation plans. Lastly, the chosen plan is executed and the concrete model is adapted. The plan is divided in two parts: (i) a model-independent part, that references only meta-level elements, and (ii) a model-specific part, that uses the functions and parameters of the concrete model.

To evaluate the performance of a model under the current environment conditions, the adaptation model considers the costs associated with each component of the T&R meta-model. Two additional costs must also be considered: the cost of staying idle due to the decision of not trusting any of the available partners and the cost of deliberation. The latter depends on the complexity of the agent's reasoning and the amount of evidence used in the evaluation of trust and reputation. The sum of these costs is used to obtain the final evaluation. After the adaptation phase, a new evaluation can be made to measure the impact of the adaptation plan.

4. APPLICATION AND EVALUATION

In order to incorporate an existing model into the meta-model, its functions, algorithms, parameters or beliefs must be mapped into the meta-model's components. These mappings allow the agent to choose the goals and plans that can be applied to the concrete model. Depending on the model, not every element in the T&R meta-model will have a correspondence in the original model. In addition to the mappings, a set of valid configurations for each component of the original model must be provided. The definition of valid configurations, together with the aforementioned mappings, result in potential adaptation plans.

To evaluate the meta-model application we provide mappings to four existing models [2]. Two of these models where chosen for experimental tests.

In the first experiment, we observe how changes in the operating costs affect the decision making component of Marsh's model [3]. In this model, a cooperation threshold is used to determine whether an agent will trust another. If trust is below this threshold, cooperation does not take place. In this experiment, it is possible to observe that agents with static cooperation threshold do not respond to changes in the operating costs and, as such, their utility is reduced as soon as the operating costs are increased. When using the meta-model, the agents reason about the effects of the changing operating costs and change the cooperation threshold accordingly. Consequently, the agent can improve its utility and avoid excessive idleness due to mistrust in a high operating cost environment without incurring in excessive loss of utility as a result of the increased risk.

In the second experiment, we observe the impact of the communication costs/information price on the process of acquiring reputation information in the FIRE model [4]. To do this, we use FIRE's information source for obtaining witness information called a referral network. The adaptation plan used in this experiment changes two parameters of the model: the branching factor (number of referrers directly contacted) and the referral length threshold (the maximum number of referrers in a referral path). In this plan, an agent first reduces the referral length and then reduces the branching factor, thus reducing the costs involved. By using the meta-model to reason about reputation information acquisition process, the agent was able to reduce the cost of information acquisition in response to an increase in communication costs, without significant loss of utility caused by changes in the reliability of the reputation evaluation.

5. CONCLUSIONS

The proposed meta-model was successfully used with two T&R models in a dynamic setting. The experimental results show the importance of the adaptation model as a complement to models used in an open MAS. Agents using the meta-model where able to adapt their models to the environmental conditions and improve the overall utility. Agents using static configurations, on the other hand, had poorer results. Further research should explore learning methods for discovering new adaptation plans and the possibility of multi-model composition.

REFERENCES

[1] I. Pinyol and J. Sabater-Mir. Computational trust and reputation models for open multi-agent systems: a review. *Artificial Intelligence Review*, July 2011.

[2] B. W. P. Hoelz and C. G. Ralha. Towards a cognitive meta-model for adaptive trust and reputation in open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 29(6):1125–1156, 2014.

[3] S. P. Marsh. *Formalising trust as a computational concept*. Doctoral thesis, University of Stirling, United Kingdom, 1994.

[4] T. D. Huynh. *Trust and reputation in open multi-agent systems*. Doctoral thesis, University of Southhampton, 2006.