



Enhancing the Adaptation of BDI Agents Using Learning Techniques

Stéphane Airiau, University of Amsterdam, The Netherlands

Lin Padgham, RMIT University, Australia

Sebastian Sardina, RMIT University, Australia

Sandip Sen, University of Tulsa, USA

ABSTRACT

Belief, Desire, and Intentions (BDI) agents are well suited for complex applications with (soft) real-time reasoning and control requirements. BDI agents are adaptive in the sense that they can quickly reason and react to asynchronous events and act accordingly. However, BDI agents lack learning capabilities to modify their behavior when failures occur frequently. We discuss the use of past experience to improve the agent's behavior. More precisely, we use past experience to improve the context conditions of the plans contained in the plan library, initially set by a BDI programmer. First, we consider a deterministic and fully observable environment and we discuss how to modify the BDI agent to prevent re-occurrence of failures, which is not a trivial task. Then, we discuss how we can use decision trees to improve the agent's behavior in a non-deterministic environment. [Article copies are available for purchase from InfoSci-on-Demand.com]

Keywords: BDI Agent-oriented Programming; Decision Trees; Learning

INCORPORATING LEARNING IN BDI AGENTS

Introduction

It is widely believed that learning is a key aspect of intelligence, as it enables adaptation to complex and changing environments. Agents developed under the Belief-Desire-Intention

(BDI) approach (Bratman, Israel, & Pollack, 1988) are capable of *simple* adaptations to their behaviors, implicitly encoded in their *plan library* (a collection of pre-defined *hierarchical plans* indexed by goals and representing the standard operations of the domain). This adaptation is due to the fact that (i) execution relies entirely on *context sensitive subgoal expansion*, and therefore, plan choices at each level of abstraction are made in response to the

current situation; and (ii) if a plan happens to fail, often because the environment has changed unexpectedly, agents “backtrack” and choose a different plan-strategy.

However, BDI-style agents are generally unable to go beyond such level of adaptation, in that they are confined to what their pre-defined plan libraries encode. As a result, they cannot significantly alter their behaviors from the ones specified during their initial deployment. In particular, these agents are not able to *learn* new behaviors (i.e., new plans) nor to learn how to choose among plans—both plans and their contexts are hard-coded. In this work, we are concerned with the latter limitation. We therefore analyze BDI-based agent designs and identify opportunities and mechanisms for performing plan context learning in typical BDI-style agents. By doing so, agents can enjoy a higher degree of adaptability by allowing them to improve their plan selection *on the basis of analysis of experiences*.

Research in machine learning can be broadly categorized into *knowledge-rich* and *knowledge-lean* techniques. Whereas some researchers have proposed and investigated learning mechanisms that incorporate and utilize significant amounts of domain knowledge (DeJong & Mooney, 1986, Ellman, 1989, Kolodner, 1993), the large majority of popular learning techniques assume very little domain knowledge and are largely data, rather than model, driven (Aha, Kibler, & Albert, 1991, Booker, Goldberg, & Holland, 1989, Kaelbling, Littman, & Moore, 1996, Krause, 1998, Quinlan, 1986, Rumelhart, Hinton, & Williams, 1986). Research in multiagent learning (Alonso, d’Inverno, Kudenko, Luck, & J.Noble, 2001, Panait & Luke, 2005, Tuyls & Nowé, 2006) has also followed this trend. This is particularly unfortunate as practical multiagent systems are meant to leverage existing domain knowledge in order to facilitate scalability, flexibility, and robustness. For most such online, real-time multiagent systems, individual agents need to quickly and effectively respond to unforeseen events as well as to gradual changes in environmental conditions. In this context, the amount of

experience and adaptation time available will be orders of magnitude less than what is assumed by offline knowledge-lean learning algorithms. As a result, techniques that take advantage of the available domain knowledge to aid and guide the learning and adaptation process are key to the development of successful agent learning approaches.

So, in the context of BDI agents, we foresee significant synergistic possibilities for combining *learning* and *reasoning* mechanisms. Whereas available domain knowledge of BDI agents can inform and direct embedded learning modules, the latter can incrementally adapt and update components of the reasoning module to “tune” the agents’ behaviors. It may well be the case that while substantial knowledge is encoded at design-time, there are additional *nuances* which can be learnt over time that will eventually yield better overall performance. In this article, then, we discuss issues and propose preliminary techniques in order to refine coarse heuristics provided by the BDI programmer at design time for the purposes of doing plan selection, that is, we focus on mechanism for “improving” the *context conditions* of existing plans.

The rest of the article is organized as follows. In the next section, we provide an overview of the relevant aspects of typical BDI-style agents. We then discuss modifications to the BDI framework so as to include mechanisms for improving plan selection by refining the context conditions of plans. We do so by relying on a number of simplifying assumptions that make the scenario an “ideal” one. After that, we outline ways in which these assumptions may be lifted. We end the article by drawing conclusions and future lines of work.

BDI AGENT-ORIENTED PROGRAMMING

The BDI (Belief-Desire-Intention) model is a popular and well-studied architecture of agency for intelligent agents situated in complex and dynamic environments. The model has its roots

16 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/article/enhancing-adaptation-bdi-agents-using/1393

Related Content

An Intelligent Model for DDoS Attack Detection and Flash Event Management

Oreoluwa Carolyn Tinubu, Adesina Simon Sodiya, Olusegun Ayodeji Ojesanmi, Emmanuel Oyeyemi Adeleke and Ahmad Alfawwaz Timehin (2022). *International Journal of Distributed Artificial Intelligence* (pp. 1-15).

www.irma-international.org/article/an-intelligent-model-for-ddos-attack-detection-and-flash-event-management/301212

Biometric Systems in Mobile Devices: A Study

Sukhdev Singhand Chander Kant (2019). *International Journal of Distributed Artificial Intelligence* (pp. 18-24).

www.irma-international.org/article/biometric-systems-in-mobile-devices/248479

Attending to Temporal Assumptions May Enrich Autonomous Agent Computer Simulations

Gus Koehler (2011). *Developments in Intelligent Agent Technologies and Multi-Agent Systems: Concepts and Applications* (pp. 1-18).

www.irma-international.org/chapter/attending-temporal-assumptions-may-enrich/49352

Locality-Sensitive Non-Linear Kalman Filter for Target Tracking

Ben-Bright Benuwa and Benjamin Ghansah (2021). *International Journal of Distributed Artificial Intelligence* (pp. 36-57).

www.irma-international.org/article/locality-sensitive-non-linear-kalman-filter-for-target-tracking/275262

Supporting Demand Supply Network Optimization with Petri Nets

Teemu Tynjala (2009). *Distributed Artificial Intelligence, Agent Technology, and Collaborative Applications* (pp. 366-383).

www.irma-international.org/chapter/supporting-demand-supply-network-optimization/8610