Chapter XI
Training Coordination Proxy Agents Using Reinforcement Learning

Myriam Abramson
Naval Research Laboratory, USA

ABSTRACT

In heterogeneous multi-agent systems, where human and non-human agents coexist, intelligent proxy agents can help smooth out fundamental differences. In this context, delegating the coordination role to proxy agents can improve the overall outcome of a task at the expense of human cognitive overload due to switching subtasks. Stability and commitment are characteristics of human teamwork, but must not prevent the detection of better opportunities. In addition, coordination proxy agents must be trained from examples as a single agent, but must interact with multiple agents. We apply machine learning techniques to the task of learning team preferences from mixed-initiative interactions and compare the outcome results of different simulated user patterns. This chapter introduces a novel approach for the adjustable autonomy of coordination proxies based on the reinforcement learning of abstract actions. In conclusion, some consequences of the symbiotic relationship that such an approach suggests are discussed.

INTRODUCTION

Advances in communication technologies has led to increased agent interactions and increased complexity in the decision-making process. To deal with this added burden, the coordination role is delegated to a proxy agent. Coordination proxy agents [Scerri et al., 2003] are personal agents that take on the coordination role on behalf of a human user (Fig. 1). While the optimization of the global task can be better achieved by the self-organization of proxy agents in dynamic environments, switching roles or teams involves preferences, such as loyalty, boredom, and persistence thresholds, in addition to interpretations that might need to be elicited from the human in the loop. For example, individual drivers differ in their tendency to switch lanes in urban traffic; truck drivers might prefer a less optimal route going through their favorite spots. This chapter addresses issues in determining when switching roles

or teams is appropriate to satisfy both the urgency of the subtask relative to the global task, the preferences of the user, and when input from the user is warranted. We hypothesize that a distinct class of agents, proxy agents, will emerge at the junction of the human and non-human worlds that will take on not only decision-making tasks such as coordination, but also the social interactive task and the adaptation task on our behalf. We envision those agents to be embedded in personal mobile devices such as cell phones and personal digital assistants and personalized through a training process.

In this chapter, we claim that through result-driven reinforcement learning, the human can train coordination proxies in a task with examples biasing the way the task is achieved with respect to the outcome of the task in a multiagent system. Similarly, in mixed-initiative planning involving goal selection, directives from the user are obtained interactively in case of plan conflict or provided a priori in the form of plan constraints. Mixed-initiative interactions in multi-agent systems provide a flexible way to harness the cognitive capabilities of the human in the loop in solving a problem while delegating more mundane tasks to the proxy agents. As in the turn-taking problem found in dialog management [Allen, 1999], the key decisions for mixed-initiative

interactions, as applied to the adjustable autonomy of proxy agents, include knowing *when* to ask for help, *when* to ask for more information, and *when* to inform the user of a decision. This chapter claims that learning user preferences is not sufficient for training coordination proxies if those preferences conflict with other agents' preferences and affect the outcome of the task. As long as preferences are inconsistent with each other as evidenced by the outcome of the task, a proxy agent must keep training and continue interacting while suggesting alternatives.

This chapter is organized as follows. A learning approach for training coordination proxies in making decisions is first introduced. We then motivate experiments in the prey/predator canonical coordination domain and present empirical results and an analysis of our evaluation. Finally, we conclude with a summary of related work and extrapolate on the consequences of such interactions. The key contribution of this work is a mixed-initiative approach based on the reinforcement learning of abstract actions and its algorithm scalable to large state space for the adjustable autonomy problem of coordination proxy agents.

Figure 1. Example of coordination proxies helping in traffic by negotiating the road
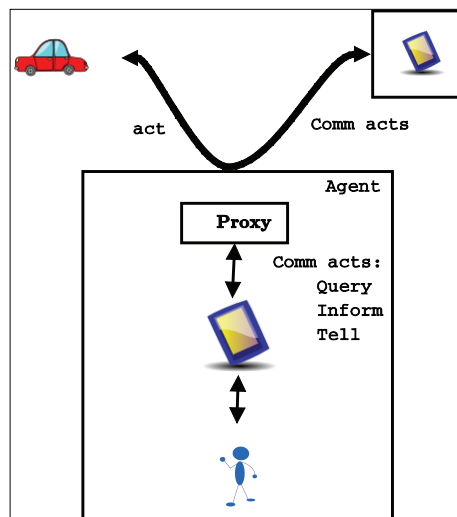
Related Content

Design of Multi Agent System for Resource Allocation and Monitoring
Manish Aroraand M. Syamala Devi (2011). *International Journal of Agent Technologies and Systems (pp. 1-10).*
www.irma-international.org/article/design-multi-agent-system-resource/52090

Enhanced Reputation Model with Forgiveness for E-Business Agents
Radu Burete, Amelia Badica, Costin Badicaand Florin Moraru (2011). *International Journal of Agent Technologies and Systems (pp. 11-26).*
www.irma-international.org/article/enhanced-reputation-model-forgiveness-business/52091

Multi-Agent Architecture for Knowledge-Driven Decision Support
Rahul Singh (2009). *Distributed Artificial Intelligence, Agent Technology, and Collaborative Applications (pp. 28-46).*
www.irma-international.org/chapter/multi-agent-architecture-knowledge-driven/8593

On Drives
Goran Trajkovski (2007). *An Imitation-Based Approach to Modeling Homogenous Agents Societies (pp. 21-29).*
www.irma-international.org/chapter/drives/5091

Evaluating Utilization of Cloud Computing for IoT Big Data Systems
Muzafar Ahmad Bhatand Amit Jain (2018). *International Journal of Distributed Artificial Intelligence (pp. 34-42).*
www.irma-international.org/article/evaluating-utilization-of-cloud-computing-for-iot-big-data-systems/238118