Chapter 12

Learning social interactions between agents

Ville Könönen, Timo Honkela, Tiina Lindh-Knuutila, Mari-Sanna Paukkeri

12.1 Introduction

One important feature of an intelligent agent is its ability to make rational decisions based on its current knowledge of the environment. If the environment of the agent is not static, i.e., there are other active entities, e.g., other agents or humans in the environment, it is crucial to model these entities for making rational decisions.

Our earlier research has been concentrated on the theoretical aspects of the modeling other agents in the reinforcement learning framework. Reinforcement learning is a learning paradigm located between supervised and unsupervised learning. In an environment, the agent takes actions and receives reward signals corresponding to the success of these actions. Correct answers are not directly provided to the agent but it learns features of the environment by continuously interacting with it.

12.2 Applications of multiagent reinforcement learning

Reinforcement learning methods have attained lots of attention in recent years. Although these methods and procedures were earlier considered to be too ambitious and to lack a firm foundation, they have been established as practical methods for solving, e.g., Markov Decision Processes (MDPs). However, the requirement for reinforcement learning methods to work is that the problem domain in which these methods are applied obeys the Markov property. Basically this means that the next state of a process depends only on the current state, not on the history. In many real-world problems this property is not fully satisfied. However, many reinforcement learning methods can still handle these situations relatively well. Especially, in the case of two or more decision makers in the same system the Markov property does not hold and more advanced methods should be used instead. A powerful tool for handling these highly non-Markov domains is the concept of Markov game. In this section we introduce two applications of multi-agent reinforcement learning, namely a dynamic pricing problem and a communication game between agents.

Dynamic pricing

A dynamic pricing scenario is a problem domain that requires planning and therefore it is an ideal testbed for reinforcement learning methods. In the problem, there are two competing brokers that sell identical products to customers and compete on the basis of price. We have modeled the problem as a Markov game and solve it by using two different learning methods. The first method utilizes modified gradient descent in the parameter space of the value function approximator and the second method uses a direct gradient of the parameterized policy function. [1]

Communication game between agents

As another problem, we consider a multiagent system, where unsupervised learning is used in the formation of agents' conceptual models. An intelligent agent usually has a goal. For studying agent based systems formally, it is useful that the goal can be expressed mathematically. Traditional approach is to define a utility function for the agent, i.e. there is a scalar value connected to each possible action measuring the fitness of the action choice for satisfying the goal of the agent. The utility function is often initially unknown and must be learned by interacting with the environment, e.g. by communicating with other agents. [2]

The agents communicate and learn through communication leading into intersubjective sharing of concepts. Communication can be modeled as a mathematical game and the structure of the game can be learnt by using reinforcement learning methods.

References

[1] V. Könönen (2006). Dynamic Pricing Based on Asymmetric Multiagent Reinforcement Learning. *International Journal of Intelligent Systems*, vol. 21, pp. 73–98.

[2] T. Honkela, V. Könönen, T. Lindh-Knuutila and M. Paukkeri (2006). Simulating processes of language emergence, communication and agent modeling. In *Proceedings of STeP-2006, the 12th Finnish Artificial Intelligence Conference*, pp. 129–132. Espoo, Finland.

12.3 Meaning negotiation using simulated naming games

The acquisition of concepts can be viewed as the process of grounding them in language use. How do concepts form, to begin with, when no language exists in the environment? The relationship between cognitive and linguistic development can be pinpointed to the question: how do conceptual emergence and the formation of names for the concepts connect? An associated important question is how an agreement on the use of words is reached in a community of agents. This process of converging towards a shared use of words is called meaning negotiation.

In [4], we studied the emergence of associations between concepts and words with help of a computer simulation and consider the hypothesis that concepts are modeled as areas in some conceptual space (see [1]). Furthermore, we utilized the self-organizing map [3] as a model for an agent's conceptual memory [2]. The feasibility of this conceptual memory model was then studied using multi-agent language games, more specifically, naming games, in which agents learn to associate words to meanings in a communicative setting. Prior to learning word-meaning-associations, a conceptual representation is also learned individually by each agent based on sensory data.

Examples of conceptual maps from simulations with six agents are shown in Fig. 12.1. They show the conceptual space of one agent after the simulation run. The self-organizing map is labeled with the words that have been used during the simulation.

Figure 12.1: An example conceptual map from the simulation. The shades of gray denote distances in the original space: The larger the distance, the darker the color. One can see that there are eight clusters on each map which corresponds to the number of prototypical colors used in the input.

The results of the experiments show clearly that when using the observational game model and the SOM-based conceptual maps (1) the agents learned to communicate successfully on the topics of the games, and (2) a shared lexicon was developed during the simulations. According our definition of successful communication, the agents are also able to communicate successfully and develop a shared lexicon based on adaptation.

References

[1] P. Gärdenfors, *Conceptual Spaces.* MIT Press, 2000.

[2] T. Honkela. Neural nets that discuss: a general model of communication based on self-organizing maps. In S. Gielen and B. Kappen, editors, *Proceedings of ICANN'93, International Conference on Artificial Neural Networks*, pages 408–411, Amsterdam, the Netherlands, September 1993. Springer-Verlag, London.

[3] T. Kohonen (2001). *Self-Organizing Maps.* Third, extended edition. Springer.

[4] T. Lindh-Knuutila, T. Honkela, and K. Lagus (2006). Simulating Meaning Negotiation Using Observational Language Games. In *Symbol Grounding and Beyond: Proceedings of the Third International Workshop on the Emergence and Evolution of Linguistic Communication*, pp. 168–179. Springer.