Machine Learning, 33, 1?? (1998)

Guest Editorial

MICHAEL HUHNS
CENTER FOR INFORMATION TECHNOLOGY
DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
UNIVERSITY OF SOUTH CAROLINA
COLUMBIA, SC 29208 USA
huhns@sc.edu

GERHARD WEISS
COMPUTER SCIENCE DEPARTMENT
TECHNICAL UNIVERSITY OF MUNICH
D-80290 MUNICH, GERMANY
weissg@informatik.tu-muenchen.de

About Multiagent Learning

This special issue is on multiagent learning, that is, on learning that relies on or even requires the interaction among several intelligent agents. An agent is commonly understood as a computational or natural entity that can be viewed as perceiving in and acting upon its environment, as being autonomous in that its behavior is at least partially determined by its own experience, and as pursuing goals or carrying out tasks (see, e.g., (Huhns & Singh, 1998) for a contemporary collection of articles on agents and multiagent systems). Multiagent learning emerged as a topic of active research in the late 1980s and early 1990s, and since then has attracted steadily increasing attention in both the multiagent systems and distributed artificial intelligence community (e.g., Bond & Gasser, 1988; Gasser & Huhns, 1989; Huhns, 1987; O'Hare & Jennings, 1996) and the machine learning community. This attention can be attributed to two primary insights:

1. There is a strong need for learning techniques and methods in the area of multiagent systems. These systems show several characteristics that make it particularly difficult to specify them correctly and completely: for instance, there is no global system control, each agent usually has just incomplete information, the information owned by different agents can be contradictory, and typically the agents are intended to operate in complex—open, large, dynamic, and unpredictable—environments. Because of these characteristics, it is obviously desirable that the agents themselves are capable of improving their own behavior, in addition to the overall system's behavior.

2. The machine learning area can profit from an extended view capturing both single-agent and multiagent learning. It is one of the primary concerns of this area to understand the principles and mechanisms of learning, whether it occurs

in computational or natural systems. Achieving such an understanding requires considering potential learners not just as "stand-alone entities" that act in isolation, but also as "social entities" that interact with one another. This obviously holds for humans and other animals as it lies in their very nature to live and act together, as well as for computing systems as they become more and more connected with each other through long-range and local-area networks.

Compared to single-agent learning, multiagent learning raises several qualitatively new issues centered around the relationship between learning and interaction. These issues can be divided into two groups:

1. *The role of interaction for learning.* Interacting agents, as they exchange information or modify the shared environment in which they are embedded, can significantly influence each other in their individual learning. Possible forms of influence are, for instance, initiation, acceleration, redirection, and prevention of another agent's learning process. Interaction makes it possible that learning by one agent can considerably change the conditions for learning with which other agents have to cope. In particular, interaction is the key to various forms of collective learning in which several agents try to achieve as a group what the individuals cannot, by sharing the load of learning and by pursuing a common learning goal on the basis of their diverse knowledge, capabilities, experience, preferences, and so forth.

2. *The role of learning for interaction.* Several dimensions of multiagent interaction can be subject to learning. These include: when to interact, with whom to interact, how to interact, and what exactly the content of interaction should be. An important pattern of multiagent interaction is coordination, among both cooperative and competitive agents. Many learning approaches to coordination are possible. For instance, agents can learn to predict the behavior of others, they can learn to detect and resolve conflicts among their planned activities, they can learn to use a common ontology, they can learn to develop shared viewpoints and assumptions, they can learn to form organizational structures (usually called teams or groups) that enable them to fulfill their design objectives, and they can learn to reconfigure their styles of coordination to respond best to environmental changes.

It is clear that these issues do not arise in single-agent contexts. There are differences in both the potential paths and the potential goals of learning in single-agent and multiagent settings, and this justifies our contention that multiagent learning is more than a mere magnification of single-agent learning.

Researchers in DAI and multiagent systems have found that knowledge representation and reasoning are different for teams of agents and for societies of agents, than they are for individual agents. A group—a team or society—might know something that no individual in the group knows. For example, a majority of the group might prefer chocolate ice cream to vanilla ice cream, but the individuals might be aware only of their own preferences.

Similarly, learning should be different for teams and societies than for individuals. The extent of a society is not fixed and is not necessarily known to any members.

Tasks and goals might not be defined or agreed upon, and measures of their success or satisfaction might also not be agreed upon. In such an environment, coordinated behavior is a challenge, but certainly requires learning, both in individual knowledge and in group knowledge. These are appropriate and still open issues for the machine learning research community.

About the Papers

This special issue brings together experimental and theoretical state-of-the-art research on multiagent learning. It includes six papers, each carefully reviewed by experts in machine learning and multiagent systems, that reflect the broad spectrum of multiagent learning and focus on different key aspects of this kind of learning.

The first paper in this issue, "Learning to improve coordinated actions in cooperative distributed problem-solving environments" by Toshiharu Sugawara and Victor Lesser, concentrates on how multiple agents can learn to identify what information can improve coordination in specific problem-solving contexts. The described work starts out from the observation that coordination is an essential technique for jointly solving problems, but that coordination strategies are not always effective and efficient in all problem domains. Sugawara and Lesser introduce and discuss an approach to learning situation-specific control rules that allow agents to identify and avoid uncoordinated situations and thus to improve the coherence of the overall distributed problem solving process. For an experimental analysis of the strengths and limitations of this learning approach, the diagnosis of local area networks was chosen as an application domain.

The next paper, "Learning coordination strategies for cooperative multiagent systems" by Fenton Ho and Mohamed Kamel, investigates how individual and collective learning can be combined to achieve coordination among multiple agents. The described work is motivated by the observation that a hand-coding of coordination strategies is very difficult and that many existing learning approaches suffer from problems with convergence, credit assignment, and complexity. A novel learning approach called multiagent probabilistic hill-climbing is introduced that addresses these problems. A basic feature of this approach is that learning occurs in two stages: in the first stage, the agents learn individually to restrict the space of potential interactions, and in the second stage the agents learn collectively to combine the results of restriction. A synthetic symbol domain and the predator-prey domain are chosen to empirically analyze this learning approach.

In "Conjectural equilibrium in multiagent learning," Michael Wellman and Junling Hu describe work centered around the question of how learning processes in multiagent systems can be generally characterized. Such a characterization would not only improve our general understanding of multiagent learning methods, but also lighten the task of designing and analyzing them—a task that is particularly challenging, because an agent's learning activities and effects can significantly change the environment for other agents. Wellman and Hu propose to use the concept of conjectural equilibrium, where the expectations of all agents are realized and each agent responds optimally to its expectations, as the central element of such a characterization. The paper presents theoretical and experimental results on the dynamics of learning in multiagent environments and on conditions for converging to conjectural equilibrium. Synthetic markets in which competitive agents interact are chosen as an application domain.

John W. Sheppard, in "Co-learning in differential games," explores multiagent learning in the context of game playing. Most work available in this context deals with games in which a single player attempts to learn a strategy that is optimal against a fixed strategy applied by its opponent. The work reported in this paper extends this view and assumes that competitive players attempt to simultaneously learn their optimal strategies. In this case, each player must be sensitive to the fact that the other player's strategy and thus the appropriateness of its own strategy varies over time. Two novel approaches to learning in competitive multiagent systems—a memory-based algorithm called MBCL and a decision tree-based algorithm called TCBL—are described and experimentally evaluated. The four games chosen for evaluation are a game of force in which two players attempt to make a falling object to land at a certain point, and three variations of the pursuit game.

The fifth paper, "Elevator group control using multiple reinforcement learning agents" by Robert H. Crites and Andrew G. Barto, presents an application of multiagent reinforcement learning to large scale dynamic optimization. As a concrete problem of practical utility, the elevator group supervisory control is chosen. In the proposed learning approach, a team of agents is used for optimization, each of which is responsible for controlling one elevator car. None of the agents is given explicit access to the actions of the others, and so cooperation has to be learned indirectly on the basis of the global reinforcement signal that is provided to the overall team. The agents employ Q-learning and use feedforward neural networks to store their action-value estimates. Experiments with two different implementations of the learning approach are reported: a parallel implementation where the agents use a central set of shared networks, and a decentralized implementation where the agents have their own independent networks. As a standard for performance comparison, existing heuristic elevator control algorithms are used.

The sixth and final paper in this issue, "Learning team strategies with multiple policy-sharing agents: A soccer case study" by Rafał Sałustowicz, Marco Wiering, and Jürgen Schmidhuber, offers an experimental comparison of two classes of reinforcement learning algorithms in multiagent contexts. The first includes learning algorithms that use state-action evaluation functions to search through the space of potential activity policies; the second includes algorithms that directly search through the policy space. As representatives of these two classes the authors choose the widely used TD-Q learning with linear neural networks and their own learning approach called Probabilistic Incremental Program Evolution. Simulated soccer serves as an application domain. What makes this domain, which has become a standard testbed in the area of multiagent systems and distributed artificial intelligence in recent years, particularly attractive, is that it includes both elements of cooperation (within a team) and competition (among teams).

A Few Concluding Remarks

In recent years several workshops and a symposium were organized that concentrated on multiagent learning:

- IJCAI-95 Workshop on Adaptation and Learning in in Multiagent Systems (Montreal, Canada, August 1995);
- ECAI-96 Workshop on Learning in Distributed Artificial Intelligence Systems (Budapest, Hungary, August 1996);
- ICMAS-96 Workshop on Learning, Interaction, and Organization in Multiagent Systems (Kyoto, Japan, December 1996);
- AAAI-97 Workshop on Multiagent Learning (Providence, USA, July 1997); and
- AAAI-96 Symposium on Adaptation, Co-evolution and Learning in Multiagent Systems (Stanford, USA, March 1996).

The papers presented at these meetings, or revised and extended versions, can be found in (Sen, 1996; Sen, 1997; Weiss, 1997; Weiss & Sen, 1996). (Other collections of papers on multiagent learning are (Sen, 1998; Weiss, 1998).) We were actively involved—as organizers, reviewers, and/or speakers—in these meetings, and can confidently say that the quality of work on multiagent learning has steadily progressed during this time. Despite the improvement, it has to be emphasized that the area of multiagent learning is still in its infancy, and that there still are many questions and problems that need to be addressed before this area will have found its defining boundaries. Apart from important progress in detail, one of the main contributions of this special issue is to help clarify the outstanding issues in this area. We think that the following working directions are of particular importance and challenge:

- Identification of general principles and concepts of multiagent learning (e.g., What are the unique requirements and conditions for multiagent learning? Are there general guidelines for designing multiagent learning algorithms?)
- Investigation of the relationships between single-agent and multiagent learning (e.g., Under what circumstances and how can a single-agent learning approach be successfully applied in multiagent environments?)
- Application of multiagent learning in complex, real-world domains. There is ongoing work in this direction (e.g., see the paper by Crites and Barto), but there are far too few application efforts compared to the increasingly important role that multiagent systems are destined to play in industrial contexts.
- Development of theoretical foundations of multiagent learning. There are, of course, mathematical results on the properties of specific approaches, but most results are of an experimental nature and there is nothing yet like a "formal theory of multiagent learning."

We hope that the reader will find this special issue both useful and interesting, and that it will foster further work on multiagent learning—perhaps along the four directions sketched above.

Acknowledgements. We are grateful to Thomas G. Dietterich for the opportunity to be the guest editors of this special issue. Our special thanks goes to the authors for making this issue possible at all, and to the reviewers for making sure that only papers of highest quality were included.

References

Bond, A. H., & Gasser, L. (1988) (Eds.). *Readings in distributed artificial intelligence*. San Mateo, CA: Morgan Kaufmann.

Gasser, L., & Huhns, M. N. (1989) (Eds.) *Distributed artificial intelligence II*. London: Pitman.

Huhns, M. N. (1987) (Ed.). *Distributed artificial intelligence*. London: Pitman.

Huhns, M. N., & Singh, M. P. (1998) (Eds.). *Readings in agents*. San Francisco, CA: Morgan Kaufmann.

O'Hare, G. M. P., & Jennings, N. R. (1996) (Eds.). *Foundations of distributed artificial intelligence*. New York: Wiley.

Sen, S. (1996) (Ed.). *Adaptation, coevolution and learning in multiagent systems* (Technical Report SS-96-01). Menlo Park, CA: AAAI Press.

Sen, S. (1997) (Ed.). *Multiagent learning* (Technical Report WS-97-03). Menlo Park, CA: AAAI Press.

Sen, S. (1998) (Guest-Ed.). Special issue on Evolution and Learning in Multiagent Systems, *International Journal of Human-Computer Studies, 48 (1)*.

Weiss, G. (1997) (Ed.). *Distributed artificial intelligence meets machine learning*. Lecture Notes in Artificial Intelligence, Vol. 1221. Berlin: Springer-Verlag.

Weiss, G. (1998) (Guest-Ed.). Special issue on Learning in Distributed Artificial Intelligence Systems, *Journal of Experimental and Theoretical Artificial Intelligence, 10 (3)*.

Weiss, G., & Sen, S. (1996) (Eds.) *Adaption and learning in multi-agent systems*. Lecture Notes in Artificial Intelligence, Vol. 1042. Berlin: Springer-Verlag.