Learning in Multi-Agent Systems*

Eduardo Alonso[1], Mark d'Inverno[2], Daniel Kudenko[3], Michael Luck[4], and Jason Noble[5]

1 Department of Computing, City University, UK
2 Cavendish School of Computer Science, University of Westminster, UK
3 Department of Computer Science, University of York, UK
4 Department of Electronics and Computer Science, University of Southampton, UK
5 School of Computing, University of Leeds, UK

August 23, 2001

1 Introduction

In recent years, multi-agent systems (MAS) have received increasing attention in the artificial intelligence community. Research in multi-agent systems involves the investigation for autonomous, rational and flexible behavior of entities such as software programs or robots, and their interaction and coordination in such diverse areas as robotics (Kitano et al., 1997), information retrieval and management (Klusch, 1999), and simulation (Gilbert and Conte, 1995). When designing agent systems, it is impossible to foresee all the potential situations an agent may encounter and specify an agent behavior optimally in advance. Agents therefore have to learn from, and adapt to, their environment, especially in a multi-agent setting.

In this panel report, we combine several different perspectives, and review some key contributing influences. The report begins with a discussion of just why learning is considered by many to be a crucial characteristic of intelligent agent systems. In the following section, the features of different learning algorithms, and their potential impact on multi-agent systems, are discussed, such as ways to achieve *multi*-agent learning, the applicability of off-line learning methods, and a discussion of the pros and cons of reactive, logic-based, and social learning methods.

2 The Relationship between Machine Learning and MAS Research

Even though machine learning (ML) has been studied extensively in the past, learning research has been mostly independent of agent research and only recently has it received more attention in connection with agents and multi-agent systems (Huhns and Weiss, 1998; Imam, 1996; Sen, 1996; Sen, 1998; Weiss and Sen, 1996; Weiss, 1997; Weiss, 1998; Weiss, 1999). This is in some ways surprising, because the ability to learn and adapt is arguably one of the most important features of intelligence (Russell and Norvig, 1995; Huhns and Singh, 1998).

Intelligence implies a certain degree of autonomy which in turn requires the ability to make independent decisions. Thus agents have to be provided with the appropriate tools to make such decisions.

*This report is the result of a panel discussion at the Third Workshop of the UK's Special Interest Group on Multi-Agent Systems.

In most dynamic domains a designer cannot possibly foresee all situations that an agent might encounter and therefore the agent needs the ability to adapt to new environments. This is especially true for multi-agent systems, where in many cases global behavior emerges rather than being pre-defined. Consequently, learning is a crucial part of autonomy and thus should be a major focus of agent and multi-agent research.

At one level, agents and multi-agent systems can be viewed as yet another application domain for machine learning systems, admittedly with its own challenges. Research taking this view is mostly reduced to applying existing (single-agent) learning algorithms more or less directly to (single) Agents in an MAS setting. Consequently, multi-agent learning is only seen as an emergent property. Even though this could be interesting from a MAS point of view, it does not seem overly challenging for ML research. Nevertheless, this is the direction most learning research for MAS has been following.

Alternatively, multi-agent systems pose the problem of distributed learning, i.e., many agents learning separately to acquire a joint hypothesis. Existing learning algorithms have been developed for single agents learning separate and independent hypotheses. Once the learning process is distributed amongst several learning agents, such learning algorithms require extensive modification, or completely new algorithms need to be developed. In distributed learning, agents need to cooperate and communicate in order to learn effectively, and these issues are being investigated extensively by MAS researchers, but to date they received only little attention in the areas of learning.

Overall, collaboration between MAS and ML researchers would be highly beneficial for both research areas, and certainly both communities can learn from each other. Fortunately, this seems to be a view that is gaining popularity, judging by the growing interest of agent researchers in ML and vice versa.

3 Aspects of Agent Learning

We are now discussing three particular issues about Multi-Agent Learning. We start off with some considerations about what the term Multi-Agent Learning means, and the difference between isolated or emergent Multi-Agent Learning and coordinated Multi-Agent Learning. In the next two subsections, different design options are presented, namely: on-line vs off line, reactive vs logic-based learning algorithms, and social learning algorithms inspired by animal learning.

3.1 Single Agent Learning vs Multi-Agent Learning

To date, most learning algorithms have been developed from a single-agent perspective. How effective can such algorithms be used in a multi-agent setting? According to Stone and Veloso (Stone and Veloso, 1998), single-agent learning focuses on how one agent improves its individual skills, irrespective of the domain in which it is embedded. That is, we cannot talk about multi-agent learning if what an agent learns does not affect nor is affected by other neighboring agents.

But can an agent situated in a multi-agent environment learn a hypothesis that does not affect or is not affected by other agents? Even if an agent is not explicitly aware of other agents, it perceives them as part of the environment and their behavior will be part of the hypothesis learned. It does seem possible to achieve coordinated group behaviour with pure single-agent learning (Sugawara and Lesser, 1993). Past research (Mundhe and Sen, 2000) has even shown that certain levels of awareness of other agents can hurt performance.

On the other hand, single agent learning might not always yield an optimal performance in multi-agent environments and there may be domains where *coordinated* multi-agent learning is a more natural metaphor and improves the effectiveness. Even though there is a difference in learning strategies depending on the level of awareness of other agents and coordination, it is an open question whether higher levels will automatically yield better performance.

3.2 On-line and Off-line Learning Methods

On-line (or incremental) learning algorithms, such as backpropagating neural networks or (in some way) reinforcement learning, have been used to compute new hypotheses incrementally as soon as a new training example becomes available. On the other hand, off-line learning methods induce a hypothesis from a set of training examples presented to the algorithm at a single time point. Obviously, on-line algorithms are better suited for multi-agent systems where agents need to update their knowledge constantly, but nevertheless it would be desirable to be able to use the large and powerful class of off-line learning algorithms as well. In order to do this, an agent needs to collect a set of training examples and then decide at some time point to compute (or re-compute) a hypothesis. The major problems to be solved are which training examples to collect (and what format they should have) and at which time point to execute the learning algorithm. The details of both decisions will mostly depend on the application domain, but one general principle could be that off-line learning is executed as soon as the current hypothesis turns out to be wrong in a number of cases above a given threshold.

One example of the application of off-line learning methods such as Inductive Logic Programming to MAS is presented in the next subsection. The specific advantage here is the ability to incorporate domain knowledge into the learning process.

4 Learning Techniques

In this section we discuss the application of reactive, logic-based, and social learning techniques to MAS.

4.1 Logic-Based Learning and Reactive Learning

In reactive systems, the overall behaviour emerges from the interaction of the component behaviours. Instead of designing protocols of coordination or providing agents with complex (BDI-like (Rao and Georgeff, 1992)) recognition models, individuals are assumed to work on value-based information (such as the distance they must keep from their neighbours) that produce social behaviour. Since internal processing is avoided, these techniques allow the agent systems to respond to the changes in their environment in a timely fashion. In Q-learning, reactive agents are given a description of the current state and have to choose the next action so as to maximize a scalar reinforcement received after each action. The task of the agent is to learn from indirect, delayed reward, to choose sequences of actions that produce the greatest cumulative rewards.

As a side effect, agents are stripped of domain knowledge that is essential for making the right decision in complex, dynamic scenarios. We cannot reduce an agent's repertoire to situation-action rules, nor simulate complex social interactions (markets, conflicts, and the like) assuming that agents do not have any domain knowledge of their environment. In order to display high-level behaviour, agents need to abstract their experience into concepts. An agent who lives without this ability must constantly invest precious energy in reacting to external stimuli. The entity that can conceptualise can turn experience into knowledge and shepherd vital resources until they are required.

Even though they are mainly off-line techniques, Knowledge-based learning techniques such as explanation-based learning (EBL) and inductive logic programming (ILP) are suitable tools for overcoming the limitations of reactive learning systems.

EBL (Carbonell, et al., 1990) has been widely used in artificial intelligence to speed-up the performance of planners. Generally speaking, the agents are concerned with improving the efficiency of the problem-solver rather than acquiring new knowledge. Obviously, problem-solvers, when presented with the same problem repeatedly, should not solve it the same way and in the same amount of time. On the contrary, it seems sensible to use general knowledge to analyse, or explain, each problem-solving instance in order to optimise future performance. This learning is not merely a way of making a program run faster, but a way of producing qualitative changes in the performance of a problem-solving system. In short, EBL extracts general rules from single examples by generating an explanation why the system succeeded or failed, and generalising it. This provides a deductive (rather than statistical) method to turn first-principles knowledge into useful, efficient, special-purpose expertise. The learned rules enable the planner to make the right choice if a similar situation arises during subsequent problem solving.

In contrast to EBL methods, ILP (Muggleton and de Raedt, 1994) computes a hypothesis based on external and initially unknown circumstances. Generally, relying exclusively on EBL-generated rules can turn out to be impractical in real-world domains in which agents work with incomplete knowledge, and thus ILP is an important addition to the system's effectiveness. ILP methods compute an inductive hypothesis with the help of training data (positive and negative examples) and background knowledge. Agents collect training examples based on executed actions and plans and their outcome. This, together with the domain knowledge base and a target predicate (e.g., success or failure) forms the basis of the inductive learning process which computes a hypothesis (i.e., a definition of the target predicate) that can be used in subsequent planning. Target predicates are given by the system designer. Once a certain number of training examples are classified incorrectly (i.e., the agent makes a certain number of mistakes in its predictions of action outcomes) a new hypothesis will be computed based on the extended training set.

Estlin (Estlin, 1998) has shown how EBL and ILP techniques can be combined in single-agent domains. The combination of EBL and ILP to produce optimal results in dynamic, complex multi-agent systems is currently being studied (Alonso and Kudenko, 1999; Kudenko and Alonso, 2000).

4.2 Social Learning

In the remainder of the paper, influences from artificial intelligence and biology are discussed with several potential learning mechanisms being outlined. These mechanisms can be seen as an alternative to logic-based approaches discussed above.

Consider a persistent multi-agent system, where new agents enter a world already populated with experienced agents. In one sense, a new agent begins with a blank slate, as it has not yet had an opportunity to learn about its environment (although it may of course be "hard-wired" with behaviours that will probably turn out to be useful). However, a new agent may not need to find out everything about the environment for itself: it may well be possible to benefit from the accumulated learning of the population of more experienced agents.

This situation could characterize highly autonomous software agents operating on the internet, for example. A new agent has not yet learned which search engine to try first, or which auction site offers the best bargains. But the situation described also matches the learning problem facing a newborn animal, especially an animal that belongs to a social species like our own. In biology, learning in multi-agent systems has been studied for some time under the heading *social learning.* There may be lessons in the biological literature for those who are interested in engineering effective multi-agent systems.

An important difference between artificial agents and animals is that in the former case we can often enforce a completely cooperative scenario, where what is good for one agent is good for them all (i.e., a common utility function). Although cooperation occurs in many animal species, the potential for conflict is never absent, because of the competition between self-replicating genes at the heart of the evolutionary process. Indeed, much of the recent work in evolutionary biology has been about conflicts of interest between individuals and how those conflicts are resolved. So, cooperative social learning may be easier to maintain, and simpler to understand, in a population of software agents than it is in a real species. However, conflicts of interest will be relevant if agents are operating in an environment with potentially malicious competitors, as will be true on the internet, for example. Social learning in such a case could involve the added complication of making sure that your "teacher" is not attempting to deceive you in order to further its own interests.

Questions that biologists ask about social learning—or any other behaviour—include "why?" and "how?" (Tinbergen, 1963). These are also referred to as questions of function and mechanism respectively. Translated into engineering terms, these questions become: when would you want to include social learning abilities in a multi-agent system, and how should you do it?

We will deal with the "why?" or "when?" question first. In recent years there has been some progress towards understanding the adaptive value of social learning. Models of cultural transmission (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1985), within-generation transmission (Laland et al., 1993; 1996), and what economists call herding behaviour (Banerjee, 1992; Bikhchandani, 1998) help to delineate the conditions under which it will be advantageous for individuals to learn from others rather than finding things out for themselves. Some of the conclusions are rather straightforward: for example, social learning is more likely to evolve when the costs of individual, trial-and-error learning are high. So, in situations where a mistake by a naive animal could mean death, perhaps through eating a poisonous plant or failing to run from a predator, we are more likely to find young animals learning from the behaviour of others. The equivalent for a software agent might be a situation where mistakes are financially costly for the agent's owner.

A more interesting finding is that social learning will be selected when the rates of change in the environment, considered either spatially or temporally, are at intermediate levels (Laland et al., 1996). The logic is as follows: in an environment that changes very slowly, hard-wired strategies (i.e., genetically transmitted information) will enable the animal to respond appropriately. If the environment changes very quickly, the animal must learn for itself based on local conditions—social learning will be inadequate because the naive animal would be trying to learn from another whose experience of the world was no longer relevant. Thus, in deciding whether or not to build a capacity for social learning into a group of software agents, we should first examine the speed at which their environment changes.

4.2.1 Mechanisms for Social Learning

Turning to the question of mechanism: there are many ways in which one agent might learn from the behaviour of another. In the social learning literature, there has long been a focus on imitation (Galef, 1988), i.e., the goal-directed copying of another's behaviour. However, as (Tomasello, 1996) points out, true imitation is a complex process that seems to involve not only perceiving and reproducing the bodily movements of another, but understanding the changes in the environment caused by the other's behaviour, and finally being able to grasp the "intentional relations" between these, i.e., knowing how and why the behaviour is supposed to bring about the goal. Much of the work on imitation has been short on specifics about the underlying mechanisms.

We will instead consider a range of simpler mechanisms that could easily be implemented in a robotic or software agents. It has long been recognized within fields like artificial life that complex global phenomena can arise from simple local rules, and this is precisely what is happening in many social learning contexts: individuals follow a simple rule (e.g., “stay close to your mother”) and, in combination with some form of learning, this gives rise to an apparently sophisticated social learning system at the group level. ¿From the point of view of building learning abilities into artificial agents, simple mechanisms have obvious advantages in terms of robustness and design costs.

Contagious behaviour is exemplified by a rule such as “If others are fleeing, flee also.” The idea is that the stimuli produced by the performance of a particular behaviour serve as triggers for others to behave in the same way. For instance, consider an animal that is “wired up” such that the characteristic sound of a conspecific moving rapidly causes it to do likewise. In a group of these animals, any stimulus that causes one of them to flee will lead to a chain reaction of rapid movements. Note that this does not involve real learning, and is merely a reactive system, but could nevertheless produce adaptive social behaviour. Possible examples of contagious behaviour include flight responses, movement in flocks or schools, and chorusing by birds and dogs (Galef, 1988). Laughing and yawning are excellent examples of contagious behaviour in humans (Provine,1996).

Stimulus enhancement (also called local enhancement) is what happens when animals obey a rule like “Follow someone older than you, and then learn from whatever happens.” For example, if you follow your parents everywhere, and your parents sometimes eat chocolate, we do not need to postulate a capacity for genuine imitation to explain the fact that you develop a liking for chocolate. Perhaps you sample chocolate pieces dropped by your parents; you then learn that chocolate-eating is good. A simple behavioural tendency—in this case, following a conspecific—combines with the capacity for learning to result in the potential transmission of acquired behaviours. (Aisner and Terkel, 1992) have shown that stimulus enhancement accounts for the acquisition of a novel feeding behaviour in black rats.

Observational learning If we add slightly more sophisticated learning abilities to stimulus enhancement, we get observational learning. The algorithm involved is approximately “Pay attention to what others are doing or experiencing, and if the results *for them* appear to be good or bad then learn from this.” (Mineka and Cook, 1988) work on fear acquisition in monkeys illustrates the idea: the authors took naive, lab-raised rhesus monkeys and allowed them to observe an experienced conspecific reacting fearfully to the presence of a snake. The observers, previously indifferent to snakes, rapidly acquired a persistent fear of them. It is easy to see that in the wild, this sort of learning could result in the transmission of acquired fears. All that needs to be assumed is that the monkeys have evolved both an innate ability to recognize the cues associated with fear on the part of a conspecific (such as grimacing and retreating), and the tendency to learn to fear a co-occurrent stimulus (i.e., the snake).

Observational learning can also exist in a simpler form: explicit evaluation of the conspecific’s experience as good or bad may be omitted. For example, Norway rats will develop a marked preference for a novel food that they smell on the breath of a conspecific (Galef, 1996). We might say that the first rat, the observer, learns that the new food is good because it observes positive consequences for the second rat, the demonstrator. That is, the observer notes that the demonstrator is still alive to tell the tale after eating a new and potentially toxic substance. It turns out, however, that the rats are not sensitive to the consequences of eating poisonous foods:

they do not learn that a food is bad if the demonstrator has become ill after eating it; in fact they develop a preference as usual. So the rats' heuristic is simply "Pay attention to what others are eating and do likewise." Noble, Todd, and Tuci (in press) simulated this phenomenon, and showed that given certain assumptions about the rats' environment (e.g., the lethality of poison and the behaviour of sick animals), their failure to evaluate the demonstrator's health is not a mistake, but is actually an adaptive strategy.

Matched-dependent behaviour Species such as rats and pigeons can readily be trained to discriminate, e.g., to press one bar when a red light is on and to press another for a green light. (Miller and Dollard, 1941) showed that this sort of learning was equally possible when the behaviour of another animal served as the discriminative stimulus; they trained rats to follow a leader left or right at a maze junction. Thus, simple reinforcement learning can result in social learning if the contingencies are right. There is no implication that the follower understands the leader's intentions, nor even that the follower is aware of the match between the leader's behaviour and its own.

Along similar lines, (Skinner, 1953) suggested that a wild pigeon could learn through trial and error that scratching in a field was likely to be rewarding (i.e., likely to result in ingesting food) if other pigeons could be seen scratching there. In fact the pigeon need not even observe others feeding: learning a correspondence between hidden food and the evidence of feeding, such as scratch marks, would amount to the same thing. The general point is that contagious behaviour may sometimes be learned.

Cross-modal matching Vocal mimicry by birds is often held to be a special case of social learning: because the original stimulus and the animal's response are in the same sensory modality, a relatively simple pattern-matching mechanism could account for the phenomenon. In contrast, copying the movements of another animal requires cross-modal matching; the observer must be able to translate the visual input associated with another's movements into appropriate motor outputs. Consider that there is no trivial link between the sight of watching someone else scratch their nose, and the experience of scratching your own nose.

None of the simple mechanisms discussed so far requires an ability to perform cross-modal matching. Even though processes like contagious behaviour or learned copying could mean that the sight of one animal doing X was a sufficient stimulus for another animal to do X, there is no suggestion of a systematic ability to copy movements. However, imagine an animal that *was* able to identify the movements of others, and map them to movements of its own muscles. If such an ability was combined with observational learning, we would get the behavioural rule "If someone else moves their head (or forelimb or tail or…) thus or so, make the same movement yourself." As with observational learning, the rule might be conditional on positive outcomes for the observed animal. Work on "mirror neurons" in monkeys (Gallese and Goldman, 1998) and humans (Iacoboni et al., 1999) is highly suggestive that, at least in primates, direct mappings may exist between movements seen and movements performed. (Meltzoff, 1996) findings on the imitative powers of very young infants also point to an innate ability to perform cross-modal matching in humans.

5 Conclusion

Learning ability is a crucial feature of intelligent agents, especially when faced with a multi-agent environment. We have presented a few of the issues involved in applying ML algorithms to MAS. We believe that there is still a lot of work to be done in the merging of the two disciplines.

Moreover, the question of using complex cognitive agents versus simple reactive ones is still a matter of debate. We have presented two major approaches representing both sides. Logic-based agents have the advantage of being able to naturally incorporate domain knowledge in the learning process, while artificial life approaches can be based on evidence from biology (e.g., nest of rats, a flock of birds, or a troop of monkeys), and much can already be achieved with such simple mechanisms.

References

R. Aisner and J. Terkel, 1992. *Cultural transmission of pine cone opening behaviour in the black rat,* Rattus rattus.

E. Alonso, and D. Kudenko, 1999. Machine learning techniques for adaptive logic-based multi-agent systems. *Proceedings of the Second Workshop of the UK Special Interest Group on Multi-Agent Systems (UKMAS-99)*, Hewlett-Packard Laboratories, Bristol.

A. V. Banerjee, 1992. A simple model of herd behavior. *Quarterly Journal of Economics*, 107:797–817.

S. Bikhchandani, D. Hirshleifer, and I. Welch, 1998. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives*, 12(3):151–170.

R. Boyd and P. J. Richerson, 1985. *Culture and the Evolutionary Process*. The University of Chicago Press, Chicago.

J. Carbonell, C. Knoblock and S. Minton. 1990. PRODIGY: An integrated architecture for planning and learning. In K. VanLehn, editor, *Architectures for Intelligence*. Hillsdale, NJ: Lawrence Erlbaum Associates.

L. L. Cavalli-Sforza and M. W. Feldman, 1981. *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton University Press, Princeton, NJ.

T. Estlin. 1998. *Using Multi-Strategy Learning to Improve Planning Efficiency and Quality*, PhD thesis, University of Texas at Austin.

B. G. Galef, Jr., 1988. Imitation in animals: History, definition, and interpretation of data from the psychological laboratory. In T. R. Zentall and B. G. Galef, Jr., editors, *Social Learning: Psychological and Biological Perspectives*, pages 3–28. Lawrence Erlbaum Associates, Hillsdale, NJ.

B. G. Galef, Jr., 1996. Social enhancement of food preferences in Norway rats: A brief review. In C. M. Heyes and B. G. Galef, Jr., editors, *Social Learning in Animals: The Roots of Culture*, pages 49–64. Academic Press, San Diego, CA.

V. Gallese and A. Goldman, 1998. Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, 2(12):493–501.

G. N. Gilbert and R. Conte, editors. 1995. *Artificial Societies: The Computer Simulation of Social Life*. UCL Press: London.

M. Huhns and M. P. Singh, editors. 1998. *Readings in Agents*. Morgan Kaufmann Publishers: San Francisco, CA.

M. Huhns and G. Weiss, editors. 1998. Special Issue on Multiagent Learning. *Machine Learning Journal*, 33(2-3).

M. Iacoboni, R. P. Woods, M. Brass, H. Bekkering, J. C. Mazziotta, and G. Rizzolatti, 1999. Cortical mechanisms of human imitation. *Science*, 286:2526–2528.

I. F. Imam, 1996. *Intelligent adaptive agents.* Papers from the 1996 AAAI Workshop, Technical Report WS-96-04, AAAI Press.

H. Kitano, Y. Kuniyoshi, I. Noda, M. Asada, H. Matsubara and E. Osawa, 1997. Robocup: A challenge problem for AI. *AI Magazine*, 18: 73-85.

M. Klusch, editor, 1999. *Intelligent Information Agents*, Springer-Verlag.

D. Kudenko and E. Alonso, 2000. Machine learning techniques for logic based multi-agent systems. In *Proceedings of The First Goddard Workshop on Formal Approaches to Agent-based Systems.* NASA Goddard Space Flight Center, Greenbelt, MD, USA. Springer-Verlag: Berlin. In press.

K. N. Laland, P. J. Richerson, and R. Boyd, 1993. Animal social learning: Toward a new theoretical approach. In P. P. Klopfer, P. P. Bateson, and N. Thomson, editors, *Perspectives in Ethology*, pages 249–276. Plenum Press, New York.

K. N. Laland, P. J. Richerson, and R. Boyd, 1996. Developing a theory of animal social learning. In C. M. Heyes and B. G. Galef, Jr., editors, *Social Learning in Animals: The Roots of Culture*, pages 129–154. Academic Press, San Diego, CA.

A. N. Meltzoff, 1996. The human infant as imitative generalist: A 20-year progress report on infant imitation with implications for cognitive psychology. In C. M. Heyes and B. G. Galef, Jr., editors, *Social Learning in Animals: The Roots of Culture*, pages 347–370. Academic Press, San Diego, CA.

N. E. Miller and J. Dollard, 1941. *Social Learning and Imitation.* Yale University Press, New Haven, CT.

S. Mineka and M. Cook, 1988. Social learning and the acquisition of snake fear in monkeys. In T. R. Zentall and B. G. Galef, Jr., editors, *Social Learning: Psychological and Biological Perspectives*, pages 51–73. Lawrence Erlbaum Associates, Hillsdale, NJ.

S. Muggleton and L. de Raedt, 1994. Inductive Logic Programming: Theory and Methods. *Journal of Logic Programming*, 19: 629-679.

M. Mundhe and S. Sen, 2000. Evaluating concurrent reinforcement learners. In *Proceedings of the Fourth International Conference on Multiagent Systems.*

J. Noble, P. M. Todd, and E. Tuci, in press. Explaining social learning of food preferences without aversions: An evolutionary simulation model of Norway rats. *Proceedings of the Royal Society of London: Biological Sciences.*

R. R. Provine, 1996. Contagious yawning and laughter. In C. M. Heyes and B. G. Galef, Jr., editors, *Social Learning in Animals: The Roots of Culture*, pages 179–208. Academic Press, San Diego, CA.

A. S. Rao and M. P. Georgeff, 1992. An abstract architecture for rational agents. In C. Rich, W. Swartout, and B. Nebel, editors, *Proceedings of Knowledge Representation and Reasoning (KR&R-92)*, pages 439-449.

S. Russell and P. Norvig, 1995. *Artificial Intelligence: A Modern Approach.* Prentice Hall.

S. Sen, 1996. *coevolution and learning in multiagent systems.* Papers from the 1996 Spring Symposium. Technical Report SS-96-01, AAAI Press.

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

B. F. Skinner, 1953. *Science and Human Behavior,* Macmillan, New York.

P. Stone and M. Veloso. 1998. Towards collaborative and adversarial learning: A case study for robotic soccer. *International Journal of Human Computer Studies*, 48.

T. Sugawara and V. Lesser, 1993. On-line learning of coordinated plans. In *Working Papers of the 12th International Workshop on Distributed Artificial Intelligence.*

N. Tinbergen, 1963. On aims and methods of ethology. *Zeitschrift für Tierpsychologie*, 20:410–433.

M. Tomasello, 1996. Do apes ape? In C. M. Heyes and B. G. Galef, Jr., editors, *Social Learning in Animals: The Roots of Culture*, pages 319–346. Academic Press, San Diego, CA.

G. Weiss and S. Sen, editors. 1996. *Adaption and learning in multiagent systems.* Lecture Notes in Artificial Intelligence, Vol. 1042. Springer-Verlag, Berlin.

G. Weiss, editor. 1997. *Distributed artificial intelligence meets machine learning.* Lectures Notes in Artificial Intelligence, Vol. 1221. Springer-Verlag, Berlin.

G. Weiss, editor, 1998. Special Issue on Learning in Distributed Artificial Intelligence Systems. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(3).

G. Weiss, 1999. Learning in Multiagent Systems. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence,* pages 559-298. The MIT Press, Cambridge, MA.