Weiß, G. Adaptation and learning in multi-agent systems: Some remarks and a bibliography, In G. Weiß and S. Sen (Eds.) (1996), *Adaption and learning in multi-agent systems* (pp. 1–21). Lecture Notes in Artificial Intelligence, Vol. 1042. Springer-Verlag.

Adaptation and Learning in Multi-Agent Systems: Some Remarks and a Bibliography

Gerhard Weiß

Institut für Informatik, Technische Universität München
D-80290 München, Germany
weissg@informatik.tu-muenchen.de

Abstract. In the last years the topic of adaptation and learning in multi-agent systems has gained increasing attention in Artificial Intelligence. This article is intended to provide a compact, introductory and motivational guide to this topic. The article consists of two sections. In the first section, "Remarks", the range and complexity of this topic is outlined by taking a general look at the concept of multi-agent systems and at the notion of adaptation and learning in these systems. This includes a description of key dimensions for classifying multi-agent systems, as well as a description of key criteria for characterizing single-agent and multi-agent learning as the two principal categories of learning in multi-agent systems. In the second section, "Bibliography", an extensive list of pointers to relevant and related work on multi-agent learning done in (Distributed) Artificial Intelligence, economics, and other disciplines is provided.

1. Remarks

Multi-Agent Systems

Multi-agent systems, that is, computational systems composed of several agents capable of mutual and environmental interaction, establish a central research and application area in Distributed Artificial Intelligence (DAI). There are four major reasons for the broad interest in multi-agent systems:

- As distributed systems they offer useful features such as parallelism, robustness and scalability, and therefore are applicable in many domains which cannot be handled by centralized systems. In particular, they are well suited for domains which require the integration of multiple sources of knowledge or activity, the resolution of interest and goal conflicts, the time-bounded processing of very large data sets, or the on-line interpretation of data arising in different geographical locations.
- The concept of multi-agent systems is in accordance with the insight gained over the past decade in disciplines like AI, psychology, and sociology that intelligence and interaction are deeply and inevitably coupled to each other. In particular, multi-agent systems realize this coupling in both directions:

on the one hand, interactivity allows the agents to increase their level of intelligence; and on the other hand, intelligence allows the agents to increase the efficiency of their interactivity.
– The study of multi-agent systems from the perspective of DAI can contribute to our understanding of natural multi-agent systems like insect societies or human teams in general, and to our understanding of complex social phenomena like collective intelligence and emergent behavior in particular.
– Today powerful computers and advanced computing networks provide a solid platform for the realization of multi-agent technology.

In the following, the concept of multi-agent systems will be described in more detail.

Differencing Aspects and their Dimensions. In the DAI literature many multi-agent systems have been described. Taking into consideration that a system always has to be considered in its environmental context in order to really understand its functionality, it can be stated that these systems differ from each other in three key aspects:

– the *environment* occupied by the multi-agent system,
– the agent-agent and agent-environment *interaction*, and
– the *agents* themselves.

For each of these differencing aspects several dimensions can be identified by which multi-agent systems can be classified. With respect to the first differencing aspect, the environment occupied by the multi-agent system, examples of such classifying dimensions (together with attributes that illustrate their spectrum of possible values) are

– the *availability* of environmental resources (ranging from restricted to ample),
– the environmental *diversity* (ranging from poor to rich),
– the environmental *uncertainty and predictability* (ranging from predictable to unpredictable), and
– the environmental *dynamics and status* (ranging from fixed to variable).

It is important to stress that it is not trivial to conclusively define the expression "environment of a multi-agent system". In particular, the widespread definition of this expression as the "sum" of the environments of the individual agents contained in the multi-agent system is problematic: because an agent's environment usually contains other agents, this definition implies that the system itself is contained in its environment. (Another problem results from the fact that an agent's environment containing other agents may be viewed as an agent on its own.) With respect to the second differencing aspect, the agent-agent and agent-environment interaction, examples of classifying dimensions are

– the *frequency* of interaction (ranging from low to high),

- the *persistence* of interaction (ranging from short-term to long-term);
- the *level* of interaction (ranging from signal passing to knowledge exchange),
- the *pattern* of interaction (ranging from unstructured to structured),
- the *variability* of interaction (ranging from fixed to changeable),
- the *type* of interaction (ranging from competitive to cooperative), and
- the *purpose* of interaction (ranging from random to goal-directed).

Finally, with respect to the third differencing aspect, the agents themselves, examples of such classifying dimensions are

- the *number of agents* involved in the multi-agent system (ranging from two upward),
- the *number of goals* an agent has (ranging from one upward),
- the *compatibility of the goals* (ranging from contradicting to complimentary),
- the *uniformity of the agents* (ranging from homogeneous to heterogeneous), and
- the *properties of the individual agents*.

Agent Properties. There has been considerable discussion and fruitful controversy on the last of these items, and the central question addressed is: *"What are the properties that let an object like a software program or an industry robot be an agent?"* Forming the intersection of the many answers which have been given to this question, one obtains something like the following "essence" of kernel properties:

- *perceptual, cognitive and effectual skills*;
- *communicative and social abilities*;
- *autonomy* (self-control).

With that, and in as far as the first two items constitute intelligence in its intuitive meaning, this "essence" implies the concise definition of an agent as an object which in some sense is intelligent and autonomous. Further properties that are often considered to be essential for agency are the following:

- *reactivity* (i.e., the ability to respond to environmental changes in reasonable time);
- *situatedness* (i.e., the ability to continuously interact with – or to be embedded in – its environment);
- *pro-activeness and deliberation* (i.e., the ability to act in a foreseeing, goal- or plan-oriented manner);
- *rationality* (i.e., the ability to always behave in a way which is suitable or even optimal for goal attainment);
- *mobility* (i.e., the ability to change the physical position);
- *introspection* (i.e., the ability to examine and self-reflect its own thoughts, ideas, plans, etc);

- *veracity* (i.e., the property of not knowingly communicating false information);
- *benevolence* (i.e., the property of always doing what is asked to do).

(Some of these terms are differently used by different authors, and the explanations provided in brackets are only intended to approximately describe their meanings.) In addition to the properties mentioned above, sometimes properties are ascribed to agents which describe their internal states. Examples of such properties or so-called *mental attitudes* are the following:

- belief, knowledge, etc, which describe *information or cognitive states*;
- intention, commitment, plan, etc, which describe *deliberative or conative states*;
- desire, goal, choice, preference, etc, which describe *motivational or affective states*.

Each of the properties listed above concerns, in one way or another, a significant aspect of agency and, with that, represents a classifying dimension for multi-agent systems.

The System-Application Assignment Problem. Clearly, it is not the attribute value of a single dimension but the combination of the attribute values of all dimensions that characterizes a multi-agent system. An understanding of the relationships between these dimensions would provide a valuable guideline for deciding which type of multi-agent system is best or at least sufficiently well suited to a given application task, and which type of application task can be best solved by a given multi-agent system. The problem of making this decision is sometimes called the (bidirectional) *multi-agent system-application assignment problem*. To solve this problem is one of the most important long-term challenges in DAI.

Topics of Current Research and Practice. There are many topics that are of relevance to the specification, implementation, handling, and assessment of multi-agent systems. These include, for instance, agent theories and architectures, communication languages, coordination mechanisms, negotiation and cooperation strategies, organization design, multi-agent planning and diagnosis, and multi-agent problem decomposition and synthesis. To discuss these topics and the specific issues raised by them would be beyond the scope and intention of this article. As a survey of the readings recommended below shows, current research and practice on agents and multi-agent systems simultanously focusses on these topics from different points of view and at different levels.

Adaptation and Learning

Adaptation and learning in multi-agent systems constitutes a further example of such a relevant topic, and it is commonly agreed by the DAI as well as the Machine Learning community that this topic deserves particular attention. As the above considerations suggest, multi-agent systems typically are of considerabe complexity with respect to both their structure and their functionality. For most application tasks, and even in environments that appear to be more or less simple, it is extremely difficult or even impossible to correctly determine the behavioral repertoire and concrete activities of a multi-agent system a priori, that is, at the time of its design and prior to its use. This would require, for instance, that it is known a priori which environmental requirements will emerge in the future, which agents will be available at the time of emergence, and how the available agents will have to interact in response to these requirements. This kind of problems resulting from the complexity of multi-agent systems can be avoided or at least reduced by endowing the agents with the ability to adapt and to learn, that is, with the ability to improve the future performance of the total system, of a part of it, or of a single agent. The rest of this section takes a closer look on the notion of adaptation and learning in multi-agent systems. In doing so, no explicit distinction is made between adaptation and learning; instead, it is assumed that "adaptation" is covered by "learning". This is in accordance with common usage, according to which the term "adaptation" is only applied to those self-modifications that enable a system to survive in a changed environment. (In its most general meaning, the term "adaptation" denotes all changes of a system so that it becomes suitable for a given situation or purpose. This meaning, however, is too broad to be of value from the viewpoint of Machine Learning.)

Categories of Learning. Learning in multi-agent systems is more than a mere magnification of learning in single-agent systems. On the one hand, learning in multi-agent systems comprises learning in single-agent systems, because an agent, although embedded in a multi-agent system, can learn in a solitary way and completely independent of the other agents. This is what can be called *single-agent or isolated learning*: learning that does not rely on the presence of multiple agents. On the other hand, learning in multi-agent systems extends learning in single-agent systems, because agents in a multi-agent system can learn in a communal way inasmuch as their learning is influenced (e.g., initiated, redirected, or made possible at all) by exchanged information, shared assumptions, commonly developed viewpoints of their environment, commonly accepted social and cultural conventions and norms which regulate and constrain their behaviors and interaction, and so forth. This is what can be called *multi-agent or interactive learning*: learning that relies on or even requires the presence of multiple agents and their interaction. Single-agent and multi-agent learning constitute the *principal categories* of learning in multi-agent systems. (There are borderline situations which make it difficult to draw clear boundaries between these two learning categories; for instance, one might think of an agent that learns about or models other agents.)

When people talk about learning in multi-agent systems, they usually think of multi-agent instead of single-agent learning. Two usages of the term "multi-agent learning" can be distinguished:

– In its stronger and *more specific meaning*, "multi-agent learning" refers only

to situations in which several agents collectively pursue a common learning goal.
– In its weaker and *less specific meaning*, "multi-agent learning" additionally refers to situations in which an agent pursues its own learning goal, but is affected in its learning by other agents, their knowledge, beliefs, intentions, and so forth.

Independent of its underlying meaning, multi-agent learning is a many-faceted activity, and therefore it is not surprising that many synonyms of this term can be found in the literature. Examples of such synonyms, each stressing another facet, are mutual learning, cooperative learning, collaborative learning, co-learning, shared learning, team learning, social learning, pluralistic learning, and organizational learning. Whereas single-agent learning has been studied in AI since decades, multi-agent learning constitutes a relatively young field of study. Compared to its age, however, this field has already reached a considerable stage of development. Multi-agent learning is the subject of the bibliography presented in the second section.

The Credit-Assignment Problem. The basic problem any learning system is confronted with is the *credit-assignment problem*, that is, the problem of properly assigning credit or blame for overall performance changes (increase and decrease) to each of the system activities that contributed to that changes. Although this problem has been traditionally considered in the context of single-agent learning, it is also existent in the context of multi-agent learning. Taking the standard AI view according to which the activities of an agent are given by the external actions carried out by it and its internal decisions implying these actions, the credit-assignment problem can be usefully decomposed into two subproblems:

– the assignment of credit or blame for an overall performance change to external actions, and
– the assignment of credit or blame for an action to the corresponding internal decisions.

The first subproblem, which might be called the *inter-agent credit-assignment problem*, is particularly difficult for multi-agent systems, because here an overall performance change may be caused by external actions of several agents. This subproblem requires that the agents answer the question "What action carried out by what agent contributed to the performance change?" The second subproblem, which might be called the *intra-agent credit-assignment problem*, is equally difficult in single-agent and multi-agent systems. This sub-problem requires that an agent answers the question "What decisions led to a contributing action?" Any approach to multi-agent learning has to attack both the inter-agent and the intra-agent subproblem in order to succeed. How difficult it is to solve these subproblems and, with that, the total credit-assignment problem, depends on the concrete learning situation.

Forms of Learning. There is a great variety in the possible forms of learning in multi-agent systems, and there are several key criteria that may be applied in order to structure this variety. Two standard examples of such criteria, which are well known in the field of ML, are the following:

- The *learning method* or strategy used by a learning entity (a single agent or several agents). The following methods are usually distinguished:
  - rote learning (i.e., direct implantation of knowledge and skills without requiring further inference or transformation from the learner);
  - learning from instruction and by advice taking (i.e., operationalization – transformation into an internal representation and integration with prior knowledge and skills – of new information like an instruction or an advice that is not directly executable by the learner);
  - learning from examples and by practice (i.e., extraction and refinement of knowledge and skills like a general concept or a standardized pattern of motion from positive and negative examples or from practical experience);
  - learning by analogy (i.e., solution-preserving transformation of knowledge and skills from a solved to a similar but unsolved problem);
  - learning by discovery (i.e., gathering new knowledge and skills by making observations, conducting experiments, and generating and testing hypotheses or theories on the basis of the observational and experimental results).

  A major difference between these methods lies in the amount of learning efforts required by them (increasing from top to bottom).
- The *learning feedback* that is available to a learning entity and that indicates the performance level achieved so far. This criterion leads to the following usual distinction:
  - supervised learning (i.e., the feedback specifies the desired activity of the learner and the objective of learning is to match this desired action as closely as possible);
  - reinforcement learning (i.e., the feedback only specifies the utility of the actual activity of the learner and the objective is to maximize this utility);
  - unsupervised learning (i.e., no explicit feedback is provided and the objective is to find out useful and desired activities on the basis of trial-and-error and self-organization processes).

  In all three cases the learning feedback is assumed to be provided by the system environment or the agents themselves. This means that the environment or an agent providing feedback acts as a "teacher" in the case of supervised learning and as a "critic" in the case of reinforcement learning; in the case of unsupervised learning, the environment and the agents just act as passive "observers".

It is important to see that different agents do not necessarily have to learn on the basis of the same learning method or the same type of learning feedback. Moreover, in the course of learning an agent may employ different learning methods and types of learning feedback. Both criteria directly or indirectly lead to the distinction between learning and teaching agents, and they show the close relationship between multi-agent learning on the one hand and teaching and tutoring on the other. Examples of other than these two standard criteria, together with a brief description of their extreme values, are the following:

- The *purpose and goal of learning*. This criterion allows to distinguish between the following two extremes (and many graduations in between them):
  - Learning that aims at an improvement with respect to one single agent, its skills and abilities.
  - Learning that aims at an improvement with respect to the agents as a unit, their coherence and coordination.

  This criterion could be refined with respect to the number and compatibility of the learning goals pursued by the agents. Generally, an agent may pursue several learning goals at the same time, and some of the learning goals pursued by the agents may be incompatible while others are complementary.
- The *decentralization* of a learning process (where a learning process consists of all activities carried out by one or more agents in order to achieve a particular learning goal). This criterion concerns the degree of distribution and parallelism, and there are two obvious extremes:
  - only one of the available agents is involved in the learning process, and the learning steps are neither distributed nor parallelized;
  - all available agents are involved, and the learning steps are "maximally" distributed and parallelized.

  Of course, the degree of dentralization may vary for different learning processes.
- An *agent's involvement* in a learning process. With respect to the importance of involvement, one can identify the following two extremes:
  - the involvement of the agent under consideration is not a necessary condition for achieving the pursued learning goal (e.g., because it can be replaced by another equivalent agent);
  - the learning goal cannot be achieved without the involvement of exactly this agent.

  Other aspects of involvement that could be applied in order to refine this criterion are its duration and intensity. It also has to be taken into consideration that an agent may be involved in several learning processes, because it may pursue several learning goals.
- The *agent-agent and agent-environment interaction* required for realizing a learning process. Two obvious extremes are the following:

- learning requires only a minimal degree of interaction;
- learning would not be possible without extensive interaction.

This criterion could be further refined with respect to the frequency, persistence, level, pattern and type of interaction.

Many combinations of different values for these criteria are possible. For instance, one might think of a small group of agents that intensively interact (by discussing, negotiating, etc) in order to understand why the overall system performance has decreased in the past, or of a large group of agents that loosly interact (by sometimes giving advices, sharing insights, etc) in order to enhance the knowledge base of one of the group members.

Challenging Research Issues. The above criteria characterize learning in multi-agent systems at the single-agent and the total-system level, and they define a large space of possible forms of multi-agent learning. Each point in this space represents a form of multi-agent learning having its specific characteristics and its specific demands on the skills and abilities of the individual agents. Research and practice in DAI and ML has just started to explore this space. Considerable progress has been made especially in the last few years, but there are still many open questions and unsolved problems. Examples of challenging issues for future research are the following:

– requirements for learning in multi-agent systems;
– principles and concepts of learning in multi-agent systems;
– models and architectures of multi-agent systems capable of learning;
– extension and transformation of single-agent learning approaches to multi-agent learning approaches;
– parallel and distributed inductive learning in multi-agent systems;
– multi-strategy and multi-perspective learning in multi-agent systems;
– learning in multi-agent systems as organizational self-design;
– theoretical analysis of learning in multi-agent systems.

In attacking these and other issues, it is likely to be very useful and inspiring to take also related work from other disciplines than (D)AI into consideration. A number of references to such work are given in the "Bibliography" section.

Selected Pointers to Related Literature

In the following, some standard pointers to the literature on DAI, agency, multi-agent systems, and single-agent learning are provided.

There is wealth of literature on DAI in general. Standard DAI books are (Bond & Gasser, 1988; Huhns, 1987; Gasser & Huhns, 1989). The first chapter of (Bond & Gasser, 1988) offers a broad overview of important aspects and problems in DAI. Traditionally, two types of DAI systems are distinguished, namely, multi-agent systems and distributed problem solving systems (see, e.g., Durfee & Rosenschein, 1994).

Those specifically interested in the various aspects of agency are referred to (Wooldridge & Jennings, 1995). The first chapter of this book, written by the book editors, provides a valuable survey of the state of the art in (D)AI research on intelligent agents. A recent agent-oriented introductory textbook on AI is presented in (Russell & Norvig, 1995).

Work on multi-agent systems can be found, e.g., in the Proceedings of the First International Conference on Multi-Agent Systems (ICMAS, 1995) as well as in the Proceedings of the European Workshops on Modelling Autonomous Agents in a Multi-Agent World (Demazeau & Müller, 1990, 1991; Werner & Demazeau, 1992; Castelfranchi & Werner, 1994; Castelfranchi & Müller, 1995).

There are many books on single-agent ML; see, e.g., the established series (Kodratoff & Michalski, 1990; Michalski, Carbonell & Mitchell, 1983, 1986; Michalski & Tecuci, 1994). A recent textbook on ML is (Langley, 1995). Actual work on ML can be found, e.g., in the Proceedings of the European and the International Conferences on Machine Learning. The credit-assignment problem of learning was first mentioned in (Minsky, 1961).

Acknowledgements

I would like to thank Daniel Hernández, Heinz-Jürgen Müller and Sandip Sen for their suggestions and comments on an earlier draft of this article.

References

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

Castelfranchi, C., & Müller, J.-P. (Eds.) (1995). *From reaction to cognition*. Lecture Notes in Artificial Intelligence, Vol. 957. Springer-Verlag.

Castelfranchi, C., & Werner, E. (Eds.) (1994). *Artificial social systems*. Lecture Notes in Artificial Intelligence, Vol. 930. Springer-Verlag.

Demazeau, Y., & Müller, J.-P. (Eds.) (1990). *Decentralized A.I.* North-Holland.

Demazeau, Y., & Müller, J.-P. (Eds.) (1991). *Decentralized A.I. 2* North-Holland.

Durfee, E.H., & Rosenschein, J.S. (1994). Distributed problem solving and multi-agent systems: Comparisons and examples. *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence* (pp. 94–104).

Gasser, L., & Huhns, M. N. (Eds.) (1989). *Distributed artificial intelligence, Vol. 2*. Pitman.

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

ICMAS (1995). *Proceedings of the First International Conference on Multiagent Systems*. AAAI Press/MIT Press.

Kodratoff, Y., & Michalski, R.S. (Eds.) (1990). *Machine learning, Vol. III*. Morgan Kaufmann.

Langley, P. (1995). *Elements of machine learning*. Morgan Kaufmann.

Michalski, R.S., Carbonell, J.G., & Mitchell, T.M. (Eds.) (1983). *Machine learning, Vol. I*. Morgan Kaufmann.

Michalski, R.S., Carbonell, J.G., & Mitchell, T.M. (Eds.) (1986). *Machine learning, Vol. II*. Morgan Kaufmann.

Michalski, R.S., & Tecuci, G. (Eds.) (1994). *Machine learning, Vol. IV*. Morgan Kaufmann.

Minsky, M. (1961). Steps towards artificial intelligence. *Proceedings of the IRE* (pp. 8–30). Reprinted in E.A. Feigenbaum & J. Feldman (Eds.) (1963), *Computers and thought* (pp. 406–450), McGraw-Hill.

Russell, S., & Norvig, P. (1995). *Artificial intelligence: A modern approach.* Prentice Hall.

Werner, E., & Demazeau, Y. (Eds.) (1992). *Decentralized A.I. 3*. North-Holland.

Wooldridge, M.J., & Jennings, N.R. (Eds.) (1995). *Intelligent agents.* Lecture Notes in Artificial Intelligence, vol. 890. Springer-Verlag.

2. Bibliography

This is a bibliography of multi-agent learning. It contains a number of references to relevant reports, articles, and books, and is intended to be an aid and service to those interested in this field.

Providing a bibliography of multi-agent learning is not without problems for three major reasons. First, multi-agent learning constitutes a relatively young but rapidly developing field of research and application. As a response to this, not only pointers to completed work, but also to work on novel ideas and of exploratory content have been included. Second, multi-agent learning constitutes a field without clear boundaries, and there are very close relationships to several other fields like single-agent learning, organizational design and adaptive systems theory. As a consequence, and apart from a few exceptions, only pointers to work that primarily deals with multi-agent learning or essential aspects of it have been included. And third, multi-agent learning constitutes a field of highly interdisciplinary nature. Therefore, not only pointers to work in (D)AI, but also to related work conducted in other disciplines have been included.

The bibliography consists of three parts. Part I contains references to work in (D)AI. (In order to avoid unnecessary redundancy, the papers in this volume are not referenced.) This part is roughly divided into two categories: "Principles, Algorithms, Applications, Tools" and "Theory". The first category contains references to work concentrating on multi-agent learning from a more practical point of view and being centered, in one or another way, around the question how learning and interaction (cooperation, communication, and so forth) in multi-agent systems are related to each other. The second category contains references to work dealing with the computational theory of team learning, which addresses the questions of efficiency and complexity of multi-agent learning from a theoretical point of view.

Part II contains references to work in economics. In this discipline multi-agent learning constitutes a traditional and well-established subject of study, where the focus of attention is on learning in organizations like business companies and state institutions. Learning in organizations, or organizational learning, is seen as a fundamental requirement for an organization's competitiveness, productivity, and innovativeness in uncertain and changing technological and market circumstances. With that, organizational learning is considered to be essential to the flexibility and sustained existence of an organization. It is likely that AI can considerably profit from the extensive knowledge about and experience with multi-agent learning that is available in economics.

Finally, part III contains a few references to work on multi-agent learning steeming from disciplines like psychology and sociology. This part is by no means complete, and the references should be just viewed as starting points for an exploration of the related literature available in these disciplines.

PART I: Work in Distributed Artificial Intelligence

Principles, Algorithms, Applications, Tools

1. Asada, M., Uchibe, E., & Hosoda, K. (1995). Agents that learn from other competitive agents. In (WS-IMLC95).
2. Boyan, J.A., & Littman, M.L. (1993). Packet routing in dynamically changing networks: A reinforcement learning approach. In J.D. Cowan, G. Tesauro & J. Alspector (Eds.), *Advances in Neural Information Processing Systems* (Vol. 6, pp. 671–678). San Fransisco: Morgan Kaufmann.
3. Brazdil, P., Gams, M., Sian, S., Torgo, L., & van de Velde, W. (1991). Learning in distributed systems and multi-agent environments. In Y. Kodratoff (Ed.), *Machine learning – EWSL-91* (pp. 412–423). Lecture Notes in Artificial Intelligence, vol. 482. Berlin: Springer-Verlag.
4. Brazdil, P., & Muggleton, S. (1991). Learning to relate terms in a multiple agent environment. In Y. Kodratoff (Ed.), *Machine Learning – EWSL-91* (pp. 424–439). Berlin: Springer-Verlag.
5. Byrne, C., & Edwards, P. (1995). Collaborating to refine knowledge. In (WS-IMLC95).
6. Chan, P.K., & Stolfo, S.J. (1993). Toward parallel and distributed learning by meta-learning. *Working Notes AAAI Workshop Know. Disc. Databases* (pp. 227–240).
7. Chan, P.K., & Stolfo, S.J. (1993). Meta-learning for multistrategy and parallel learning. *Proceedings of the Second International Workshop on Multistrategy Learning* (pp. 150–165).
8. Chan, P.K., & Stolfo, S.J. (1993). Toward multistrategy parallel and distributed learning in sequence analysis. *Proceedings of the First International Conference on Intelligent Systems for Molecular Biology* (pp. 65–73).
9. Chan, P.K., & Stolfo, S.J. (1993). Experiments on multistrategy learning by meta-learning. *Proceedings of the Second International Conference on Inform. Know. Management* (pp. 314–323).
10. Clouse, J.A. (1995). Learning from an automated training agent. In (WS-IMLC95).

11. Davies, W., & Edwards, P. (1995). Distributed learning: An agent-based approach to data-mining. In (WS-IMLC95).
12. Dorigo, M., & Gambardella, L.M. (1995). *Ant-Q. A reinforcement learning approach to combinatorial optimization*. Technical Report 95-01. IRIDIA, Université Libre de Bruxelles.
13. Dorigo, M., Maniezzo, V., & Colorni, A. (1996). The ant system: Optimization by a colony of cooperating agents. Appears in *IEEE Transactions on Systems, Man, and Cybernetics*, 26(2).
14. Dowell, M.L. (1995). *Learning in multiagent systems*. Dissertation. Department of Electrical and Computer Engineering, University of South Carolina.
15. Dowell, M.L., & Bonnell, R.D. (1991). Learning for distributed artificial intelligence. *Proceedings of the Twenty-Third Southeastern Symposium on System Theory* (pp. 218–221).
16. Dowell, M.L., & Stephens, L.M. (1994). MAGE: Additions to the AGE algorithm for learning in multi-agent systems. *Proceedings of the Second International Working Conference on Cooperating Knowledge Based Systems (CKBS94)*.
17. Edwards, P., & Davies, W. (1993). A heterogeneous multi-agent learning system. In S.M. Deen (Ed.), *Proceedings of the Special Interest Group on Co-operating Knowledge Based Systems* (pp. 163–184).
18. Findler, N.V. (1991). Distributed control of collaborating and learning expert systems for street traffic signals. In Lewis & Stephanon (Eds.), *IFAC Distributed Intelligence Systems* (pp. 125-130). Pergamon Press.
19. Gil, Y. (1995). Acquiring knowledge from users in a reflective architecture. In (WS-IMLC95).
20. Grosof, B. (1995). Conflict resolution in advice taking and instruction for learning agents. In (WS-IMLC95).
21. Huhns, M.N., Mukhopadhyay, U., Stephens, L.M., & Bonnell, R.D. (1987). DAI for document retrieval: The MINDS project. In M.N. Huhns (Ed.), *Distributed Artificial Intelligence* (pp. 249–283). Pitman.
22. Humphrys, M. (1995). *W-learning: Competition among selfish Q-learners*. Technical Report no. 362. Computer Laboratory, University of Cambridge.
23. Kinney, M., & Tsatsoulis, C. (1993). Learning communication strategies in distributed agent environments. Working Paper WP-93-4. Intelligent Design Laboratory, University of Kansas.
24. Littman, M.L. (1994). Markov games as a framework for multi-agent reinforcement learning. *Proceedings of the 1994 International Conference on Machine Learning* (pp. 157–163).
25. Littman, M.L., & Boyan, J.A. (1993). *A distributed reinforcement learning scheme for network routing*. Report CMU-CS-93-165. School of Computer Science, Carnegie Mellon University.
26. Maes, P. (1994). Social interface agents: Acquiring competence by learning from users and other agents. In O. Etzioni (Ed.), *Working Notes of the 1994 AAAI Spring Symposium on Software Agents*.
27. Markey, K.L. (1993). Efficient learning of multiple degree-of-freedom control

problems with quasi-independent Q-agents. *Proceedings of the 1993 Connectionist Models Summer School*. NJ: Lawrence Erlbaum Associates, Inc.

28. Matarić, M.J. (1994). Learning to behave socially. *Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior – From animals to animats* (pp. 453–462).
29. Michalski, R., & Tecuci, G. (1995). *Machine learning. A multistrategy approach*. San Francisco, CA: Morgan Kaufmann.
30. Nagendra Prasad, M.V., Lander, S., & Lesser, V.R. (1995). *Experiences with a Multi-agent Design System*. Technical Report. Department of Computer Science, University of Massachusetts.
31. Nagendra Prasad, M.V., Lesser, V., & Lander, S. (1995). *Learning organizational roles in a heterogeneous multi-agent system*. Technical Report TR-95-35. Department of Computer Science, University of Massachusetts.
32. Nagendra Prasad, M.V., Lesser, V., & Lander, S. (1995). Learning experiments in a heterogeneous multi-agent system. In *Working Notes of the IJCAI95 Workshop on Adaptation and Learning in Multiagent Systems* (pp. 59–64).
33. Ohko, T., Hiraki, K., & Anzai, Y. (1995). LEMMING: A learning system for multi-robot environments. *Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots vand Systems* (Vol. 2, pp. 1131–1146).
34. Parker, L. (1993). Adaptive action selection for cooperative agent teams. *Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 442–450).
35. Parker, L. (1993). Learning in cooperative robot teams. Paper presented at *IJCAI93 Workshop on Dynamically Interacting Robots*.
36. Payne, T.R., Edwards, P., & Green, C.L. (1995). Experience with rule induction and k-nearest neighbour methods for interface agents that learn. In (WS-IMLC95).
37. Pearson, D., & Huffman, S. (1995). Combining learning from instruction with recovery from incorrect knowledge. In (WS-IMLC95).
38. Provost, F.J. (1995). Scaling up inductive learning with massive parallelism. *Machine Learning*.
39. Provost, F.J., & Hennessy, D. (1994). Distributed machine learning: Scaling up with coarse-grained parallelism. *Proceedings of the Second International Conference on Intelligent Systems for Molecular Biology*.
40. Sandholm, T.W., & Crites, R.H. (1995). Multiagent reinforcement learning in the iterated prisoner's dilemma. *Biosystems*, Special Issue on the prisoner's dilemma.
41. Schaerf, A., Shoham, Y., & Tennenholtz, M. (1995). Adaptive load balancing: A study in multi-agent-learning. *Journal of Artificial Intelligence Research*, 2, 475–500.
42. Sekaran, M., & Sen, S. (1994). Multi-agent learning in non-cooperative domains. *Proceedings of the 12th National Conference on Artificial Intelligence* (Vol. 2, pp. 1489). Menlo Park, CA: AAAI Press/MIT Press.
43. Sen, S., Sekaran, M., & Hale, J. (1994). Learning to coordinate without shar-

ing information. *Proceedings of the 12th National Conference on Artificial Intelligence* (Vol. 1, pp. 426–431). Menlo Park, CA: AAAI Press/MIT Press.

44. Shavlik, J., & Maclin, R. (1995). Learning from instruction and experience in competitive situations. In (WS-IMLC95).
45. Shoham, Y., & Tennenholtz, M. (1992). Emergent conventions in multi-agent systems: initial experimental results and observations. *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning* (pp. 225–231).
46. Shoham, Y., & Tennenholtz, M. (1994). *Co-learning and the evolution of social activity*. Technical Report STAN-CS-TR-94-1511. Department of Computer Science, Stanford University.
47. Shoham, Y., & Tennenholtz, M. (1995). Social Laws for Artificial Agent Societies: Off-line Design. Appears in *Artificial Intelligence*, 73.
48. Sian, S.S. (1990). The role of cooperation in multi-agent learning. *Proceedings of the International Working Conference on Cooperating Knowledge Based Systems.* (pp. 164–177). Springer-Verlag.
49. Sian, S.S. (1991). Adaptation based on cooperative learning in multi-agent systems. In Y. Demazeau, & J.-P. Müller (Eds.), *Decentralized A.I. 2* (pp. 257–272). Amsterdam: North-Holland.
50. Sian, S.S. (1991). Extending learning to multiple agents: issues and a model for multi-agent machine learning (MA-ML). In Y. Kodratoff (Ed.), *Machine learning – EWSL-91* (pp. 440–456). Berlin: Springer-Verlag.
51. Sikora, R., & Shaw, M.J. (1990). *A double-layered learning approach to acquiring rules for financial classification*. Faculty Working Paper No. 90–1693. College of Commerce and Business Administration, University of Illinois at Urbana-Champaign.
52. Sikora, R., & Shaw, M.J. (1991). *A distributed problem-solving approach to inductive learning*. Faculty Working Paper 91-0109. College of Commerce and Business Administration, University of Illinois at Urbana-Champaign, Illinois.
53. Tan, M. (1993). Multi-agent reinforcement learning: Independent vs. cooperative agents. *Proceedings of the Tenth International Conference on Machine Learning* (pp. 330–337).
54. Tennenholtz, M. (1995). On Computational Social Laws for Dynamic Non-Homogeneous Social Structures. Appears in *Journal of Experimental and Theoretical Artificial Intelligence*.
55. Weiß, G. (1993). Collective learning and action coordination. *Proceedings of the 13th International Conference on Distributed Computing Systems* (pp. 203–209).
56. Weiß, G. (1993). Learning to coordinate actions in multi-agent systems. *Proceedings of the 13th International Conference on Artificial Intelligence* (Vol. 1, pp. 311–316).
57. Weiß, G. (1993). Action selection and learning in multi-agent environments. *Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 502–510).

58. Weiß, G. (1993). Lernen und Aktionskoordinierung in Mehragentensystemen. In J. Müller (Ed.), *Verteilte Künstliche Intelligenz – Methoden und Anwendungen* (pp. 122–132). Mannheim: BI Verlag.
59. Weiß, G. (1994). *Some studies in distributed machine learning and organizational design*. Techn. Rep. FKI-189-94. Institut für Informatik, TU München.
60. Weiß, G. (1995). Distributed reinforcement learning. *Robotics and Autonomous Systems*, 15, 135–142.
61. Weiß, G. (1995). *Distributed machine learning*. Sankt Augustin: Infix-Verlag.
62. WS-IMLC95 (1995). Workshop "Agents that learn from other agents" held at the 1995 International Machine Learning Conference. (Proceedings are ALSO available at http://www.cs.wisc.edu/~shavlik/ml95w1/pubs.html.)

Theory

63. Daley, R.P., Kalyanasundaram, B., & Velauthapillai, M. (1992). Breaking the probability 1/2 barrier in fin-type learning. *Proceedings of the Fifth Annual Workshop on Computational Learning Theory* (pp. 203–217), Pittsburgh, Pennsylvania. ACM Press.
64. Daley, R.P., Kalyanasundaram, B., & Velauthapillai, M. (1992). The power of probabilism in popperian finite learning. *Proceedings of the Third International Workshop on Analogical and Inductive Inference* (pp. 151–169). Dagstuhl Castle, Germany.
65. Daley, R.P., Kalyanasundaram, B., & Velauthapillai, M. (1993). Capabilities of fallible finite learning. *Proceedings of the Sixth Annual Conference on Computational Learning Theory* (pp. 199–208), Santa Cruz, CA. ACM Press.
66. Daley, R.P., Pitt, L., Velauthapillai, M., Will, T. (1991). Relations between probabilistic and team one-shot learners. In L. Valiant & M. Warmuth (Eds.), *Proceedings of the Workshop on Computational Learning Theory* (pp. 228–239). Morgan Kaufmann.
67. Jain, S., & Sharma, A. (1990). Finite learning by a team. In M. Fulk & J. Case (Eds.), *Proceedings of the Third Annual Workshop on Computational Learning Theory* (pp. 163–177). Morgan Kaufmann.
68. Jain, S., & Sharma, A. (1990). Language learning by a team. In M.S. Paterson (Ed.), *Proceedings of the 17th International Colloquium on Automata, Languages and Programming* (pp. 153–166). Springer-Verlag.
69. Jain, S., & Sharma, A. (1993). *Computational limits on team identification of languages*. Technical Report 9301. School of Computer Science and Engineering, University of New South Wales.
70. Jain, S., & Sharma, A. (1993). Probability is more powerful than team for language identification. *Proceedings of the Sixth Annual Conference on Computational Learning Theory* (pp. 192–198). ACM Press.
71. Jain, S., & Sharma, A. (1994). *On aggregation teams of learning machines*. SCS&E Report no. 9405. School of Computer Science and Engineering, University of New South Wales.

72. Jain, S., & Sharma, A. (1995). Team learning of formal languages. In (WS-IMLC95).
73. Jain, S., Sharma, A., & Velauthapillai, M. (1994). *Finite identification of functions by teams with success ratio 1/2 and above. Journal of Computer and System Sciences.*
74. Pitt, L. (1984). A characterization of probabilistic inference. *Proceedings of the 25th Symposium on the Foundations of Computer Science.*
75. Pitt, L. (1989). Probabilistic inductive inference. *Journal of the ACM*, 36, 383–433.
76. Pitt, L., & Smith, C. (1988). Probability and plurality for aggregations of learning machines. *Information and Computation*, 77, 77–92.
77. Smith, C. (1982). The power of pluralism for automatic program synthesis. *Journal of the ACM*, 29, 1144–1165.
78. Velauthapillai, M. (1989). Inductive inference with bounded number of mind changes. *Proceedings of the Workshop on Computational Learning Theory* (pp. 200–213).

PART II: Work in Economics

79. Adler, P. (1990). Shared learning. *Management Science*, 36/8, 938–957.
80. Argyris, C. (1982). *Reasoning, learning and action: individual and organizational.* San Fransisco: Jossey-Bass.
81. Argyris, C. (1993). *On organizational learning.* Cambridge, Mass: Blackwell.
82. Argyris, C., & Schön, D.A. (1978). *Organizational learning.* Reading, MA: Addison-Wesley.
83. Arrow, K. (1962). The implications of learning by doing. *Review of Economic Studies*, 29, 166-170.
84. Bennis, W., & Nanus, B. (1985). Organizational learning: The management of the collective self. *New Management*, 3, 6–13.
85. Brown, J.S., & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*, 2(1), 40–57.
86. Cangelosi, V.E., & Dill, W.R. (1965). Organizational learning: Observations toward a theory. *Administrative Science Quaterly*, 10, 175–203.
87. Cohen, M.D. (1986). Artificial intelligence and the dynamic performance of organizational designs. In J.G. March & R. Weissinger-Baylon (eds.), *Ambiguity and command* (pp. 53–71). Marshfield, Mass.: Pitman.
88. Cohen, M.D. (1991). Individual learning and organizational routine: Emerging connections. *Organization Science*, 2(1), 135–139.
89. Cohen, M.D. (1992). When can two heads learn better than one? Results from a computer model of organizational learning. In M. Masuch & M. Warglien (Eds.), *Artificial intelligence in organization and management theory* (pp. 175–188), Amsterdam: North-Holland.
90. Cohen, W., & Levinthal, D. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quaterly*, 35, 128–152.

91. Cohen, M.D., March, J.G., & Olsen, J.P. (1972). A garbage can model of organizational choice. *Administrative Science Quaterly*, 17(1), 1–25.
92. Cohen, M.D., & Sproull, L.S. (1991) (Eds.). *Organization Science*, 2(1) [Special Issue on Organizational Learning].
93. Daft, R.L., & Huber, G.P. (1987). How organizations learn: A communication framework. *Research in the Sociology of Organizations*, 5, 1–36.
94. Derry, D. (1983). Decision-making, problem-solving, and organizational learning. *Omega*, 11, 321–328.
95. Dixon, N.M. (1990). Action learning, action science and learning new skills. *Industrial & Commercial Training*, 22(4), 1–16.
96. Dixon, N.M. (1992). Organizational learning: A review of the literature with implications for HRD professionals. *Human Resource Development Quaterly*, 3(1), Spring, 29–49.
97. Dodgson, M. (1993). Organizational learning: A review of some literatures. *Organization Studies*, 14/4, 375–394.
98. Duncan, R., & Weiss, A. (1979). Organizational learning: Implications for organizational design. In B.M. Staw (Ed.), *Research in organizational behavior* (Vol. 1, pp. 75–123). Greenwich, Conn.: JAI Press.
99. Easterby-Smith, M. (1990). Creating a learning organisation. *Personal Review*, 19(5), 24–28.
100. Epple, D., Argote, L., & Devadas, R. (1991). Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing. *Organization Science*, 2(1), 58–70.
101. Fiol, M., & Lyles, M.A. (1993). Organizational learning. *Academy of Management Review*, 10, 803–813.
102. Friedlander, F. (1983). Patterns of individual and organizational learning. In Shrivastava and Associates (Eds.), *The executive mind: New insights on managerial thought and action* (pp. 192–220). San Fransisco: Jossey-Bass.
103. Garratt, B. (1990). *Creating a learning organization: A guide to leadership, learning and development*.
104. Garratt, B., & Burgoyne, J.G. (1987). *The learning organization.* London: Fontana/Collins.
105. Grantham, C. (1994). The learning organization. In S.A. Katsikides (Ed.), *Informatics, organization and society* (pp. 228–250). Wien: Oldenbourg.
106. Hall, D.T., & Fukami, C.V. (1979). Organization design and adult learning. In B.M. Staw (Ed.), *Research in organizational behavior* (Vol. 1, pp. 125–167). Greenwich, Conn.: JAI Press.
107. Hedberg, B. (1981). How organizations learn and unlearn. In P.C. Nystrom & W.H. Starbuck (Eds.), *Handbook of organizational design* (Vol. 1, pp. 1–27). New York: Oxford University Press.
108. Herriott, S. R., Levinthal, D., & March, J. G. (1985). Learning from experience in organizations. *The American Economic Review*, 75(2), 298–302.
109. Huber, G.P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2(1), February, 88–115.

110. Hutchins, E. (1991). Organizing work by adaptation. *Organization Science*, 2(1), 14–39.
111. Jelinek, M. (1979). *Institutionalizing innovations: A study of organizational learning systems*. New York: Praeger.
112. Kim, D. (1993). The link between individual and organizational learning. *Sloan Management Review*, Fall, 37–50.
113. Lant, T. (1994). Computer simulations of organizations as experiential learning systems: Implications for organization theory. In K.M. Carley & M.J. Prietula (Eds.), *Computational organization theory*, Lawrence Erlbaum Associates, Inc.
114. Lant, T., & Mezias, S. (1990). Managing discontinuous change: a simulation study of organizational learning and entrepreneurship. *Strategic Management Journal*, 11, 147–179.
115. Lant, T., & Mezias, S. (1990). An organizational learning model of convergence and reorientation. *Strategic Management Journal*, 11, 147–179.
116. Levinthal, D.A. (1991). Organizational adaptation and environmental selection –Interrelated processes of change. *Organization Science*, 2(1), 140–145.
117. Levitt, B., & March, J.G. (1988). Organizational learning. *Annual Review of Sociology*, 14, 319–340.
118. Lounamaa, P.H., & March, J.G. (1987). Adaptive coordination of a learning team. *Management Science*, 33, 107–123.
119. Lundberg, C. (1989). On organizational learning: Implications and opportunities for expanding organizational development. *Research in Organizational Change and Development*, 3(6), 126–182.
120. March, J.G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.
121. March, J.G., & Olsen, J.P. (1975). The uncertainty of the past: Organizational learning under ambiguity. *European Journal of Political Research*, 3, 147–171.
122. March, J.G., Sproull, L.S., & Tamuz, M. (1991). Learning from samples of one or fewer. *Organization Science*, 2(1), 1–13.
123. Marsick, V. (Ed.) (1987). *Learning in the workplace*. New York: Croom Helm.
124. Masuch, M., & Lapotin, P. (1989). Beyond garbage cans: An AI model of organizational choice. *Administrative Science Quaterly*, 34(1), 38–67.
125. Miles, R.H., & Randolph, W.A. (1980). Influence of organizational learning styles on early development. In J.R. Kimberly & R.H. Miles (Eds.), *The organizational life cycle: Issues in the creation, transformation, and decline of organizations* (pp. 44–82). San Francisco: Jossey-Bass.
126. Mills, D.Q., & Friesen, B. (1992). The learning organization. *European Management Journal*, 10(2), 146–156.
127. Mody, A. (1990). *Learning through allianes*. Washington: The World Bank.
128. Normann, R. (1985). Developing capabilities for organizational learning. In J.M. Pennings & Associates (Eds.), *Organizational strategy and change: New views on formulating and implementing strategic decisions* (pp. 217–248).

San Francisco: Jossey-Bass.

129. Nystrom, P.C., & Starbuck, W.H. (1984). To avoid organizational crises, unlearn. *Organizational Dynamics*, 53–65.

130. Pautzke, G. (1989). *Die Evolution der organisatorischen Wissensbasis: Bausteine zu einer Theorie des organisatorischen Lernens*. Herrsching: Verlag Barbara Kirsch.

131. Pedler, M., Boydell, T., & Burgoyne, J. (1989). Towards the learning company. *Management Education and Development*, 20, 1–8.

132. Perelman, L. (1984). *The learning enterprise: Adult learning, human capital and economic development*. Washington, DC: Council of State Planning Agencies.

133. Probst, G.J.B., & Büchel, B.S.T. (1994). *Organisationales Lernen. Wettbewerbsvorteil der Zukunft*. Wiesbaden: Gabler.

134. Pucik, V. (1988). Strategic alliances, organizational learning, and competitive advantage: The HRD agenda. *Human Resource Management*, 27(1), 77–93.

135. Reinhardt, R. (1993). Das Modell organisationaler Lernfähigkeit und die Gestaltung lernfähiger Organisationen. Frankfurt am Main: Peter Lang.

136. Reber, G. (1992). Lernen, organisationales. In E. Frese (Ed.), Handwörterbuch der Organisation (pp. 1240–1255).
Revans, R.W. (1980). *Action learning: New techniques for management*. London: Blond & Briggs.

137. Sattelberger, T. (Ed.) (1994). *Die lernende Organisation. Konzepte für eine neue Qualität der Unternehmensführung*. Wiesbaden: Gabler.

138. Schein, E.H. (1993). How can organizations learn faster? The challenge of entering the green room. *Sloan Management Review*, Winter, 85–92.

139. Senge, P. (1990). *The Fifth Discipline: The art and practice of the learning organization*. New York: Doubleday.

140. Simon, H.A. (1991). Bounded rationality and organizational learning. *Organization Science*, 2(1), February, 125–134.

141. Shrivastava, P. (1983). A typology of organizational learning systems. *Journal of Management Studies*, 20, 7–28.

142. Starbuck, W.H., & Dutton, J.M. (1973). Designing adaptive organizations. *Journal of Business Policy*, 3, 21–28.

143. Stata, R. (1989). Organizational learning – The key to management innovation. *Sloan Management Review*, 30(3), 63–74.

144. Warglien, M. (1992). Exit, voice, and learning: Adaptive behavior and competition in a hotelling world. In M. Masuch & M. Warglien (Eds.), *Artificial intelligence in organization and management theory* (pp. 189–214), Amsterdam: North-Holland.

145. Weick, K.E. (1991). The nontraditional quality of organizational learning. *Organization Science*, 2(1), 116–124.

146. Wolff, R. (1982). *Der Prozeß des Organisierens: Zu einer Theorie des organisationellen Lernens*. Spardorf: Wilfer.

147. Yelle, L.E. (1990). The learning curve: Historical review and comprehensive survey. *Decision Sciences*, 10, 302–328.

PART III: Others

148. Bandura, A. (1977). *Social learning theory*. Englewood Cliffs, NJ: Prentice-Hall,
149. Dillenbourg, P., Mendelsohn, P., & Schneider, D. (1994). The distribution of pedagogical roles in a multi-agent learning environment. In R. Lewis & P. Mendelsohn (Eds.), *Lessons from Learning* (pp. 199–216). Amsterdam: North-Holland.
150. Dillenbourg, P., Baker, M., Blaye, A., & O'Malley, C. (1995, to appear). The evolution of research on collaborative learning. In H. Spada & P. Reimann (Eds.), *Learning in Humans and Machines*.
151. Laughlin, P.R. (1988). Collective induction: group performance, social combination processes, and mutual majority and minority influence. *Journal of Personality and Social Psychology*, 54(2), 254–267.
152. Mandl, H., & Renkl, A. (1992). A plea for "more local" theories of cooperative learning. *Learning and Instruction*, 2, 281–285.
153. Roschelle, J. (1992). Learning by collaboration: Convergent conceptual change. *Journal of the Learning Sciences*, 2, 235–276.
154. Slavin, R.E. (1983). *Cooperative learning*. New York: Longman.