Learning in Distributed Systems and Multi-Agent Environments

P. Brazdil[1] M. Gams[2] S. Sian[3]
L.Torgo[1] W. van de Velde[4] *

[1] LIACC-CIUP, Rua Campo Alegre, 823, 4100 Porto, Portugal
E-mail: pbrazdil@nccup.ctt.pt (later pbrazdil@liacc.up.pt)

[2] Jozef Stefan Institute, Jamova 39, 61000 Ljubljana, Yugoslavia
E-mail: mezi@ijs.ac.mail.yu

[3] Imperial College, Department of Computing, 180 Queen's Gate
London SW7 2BZ, UK. E-mail: sss@doc.ic.ac.uk

[4] Vrije Universitaet, AI Lab., Pleinlaan 2, B-1050 Brussels, Belgium
E-mail: walter@arti.vub.ac.be

Abstract

The paper begins with the discussion on why we should be concerned with machine learning in the context of distributed AI. The rest of the paper is dedicated to various problems of multi-agent learning. First, a common framework for comparing different existing systems is presented. It is pointed out that it is useful to distinguish *when* the individual agents communicate. Some systems communicate during the learning phase, others during the problem solving phase, for example. It is also important to consider *how*, that is in what language, the communication is established. The paper analyses several systems in this framework. Particular attention is paid to previous work done by the authors in this area. The paper covers use of redundant knowledge, knowledge integration, evaluation of hypothesis by a community of agents and resolution of language differences between agents.

Keywords: learning in distributed systems, multi-agent learning, evaluation of hypotheses, knowledge integration, use of redundant knowledge, resolution of language differences.

(*) This paper has been prepared with a specific purpose in mind - to provide a basis for the *Panel on Learning in Distributed Systems.* It was edited by P.Brazdil on the basis of the individual contributions received and/or papers made available. The names of the authors are shown in the alphabetic order.

1. Introduction

As Sian (1991) has pointed out many real world problems are best modelled using a set of cooperating intelligent systems (agents). There are many reasons that we could give to justify our position. First, our society consists of many interacting entities and so if we are interested to model some aspects of our society, our model needs to be structured. Also, as data often originates at different physical locations, centralized solutions are often inapplicable or inconvenient. Recent work in the field of *Distributed Artificial Intelligence* (DAI) and multi-agent systems (Huhns, 1987; Bond and Gasser, 1988; Durfee et al., 1989; Demazeau et al., 1991, etc.) has addressed the issues of organization, coordination and cooperation. The problems of multi-agent learning has, however, been largely ignored. One purpose of this paper is to address some issues that arise when studying ML in multi-agent systems.

Two rather different questions can be formulated in this context. First, how can multi-agent systems benefit from machine learning. Second, how can machine learning benefit from considering multi-agent set-up. As multi-agent systems are by nature complex, machine learning techniques may be the only way to achieve a robust and versatile system. The advantages of ML cannot be taken for granted, but rather have to be demonstrated in terms of its effects on cost, time, resources and product quality. One may envisage advantages defined in terms of ease of programming, maintenance, scope of application, efficiency and coordination of activity.

One may wonder why the researchers in machine learning should venture into an area so difficult as distributed AI. We believe that multi-agent learning will touch upon some of the fundamental issues of intelligence and learning that can be only understood in this context. Although communication seems to play an important role in human learning, so far this has not been studied much in ML.

Studying multi-agent learning may help us to design systems that are faster, thanks to the possibility of parallelism. Furthermore, the systems may become more robust when compared to single-agent systems. As has been shown by various authors (e.g. Gams, 1989; Buntine, 1989; Brazdil and Torgo, 1990 etc.) cross checking of results between different methods provides more reliable results.

The study of multi-agent learning poses new questions that need to be answered. For example, when should the individual systems cooperate and how. The purpose of this paper is to discuss several different approaches that have been taken. This discussion will not attempt to be exhaustive, but rather concentrate mainly on the work done in this area by the authors of this paper. However, an attempt will be made to present this work in a unified perspective and suggest directions for further work.

The paper is organized as follows. In Section 2 we shall briefly discuss autonomous agent learning. Section 3 will be dedicated to multi-agent learning. It will describe certain criteria that we can use when comparing different systems. This section describes several existing systems and approaches and is mainly oriented towards some earlier work done by the authors. The last section will discuss new horizons and future work.

2. Autonomous Agent Learning

Multi-agent learning could be seen as an extension of autonomous agent learning. But could the study of multi-agent learning really benefit from the results that have been achieved in autonomous learning?

The application of symbolic AI to robotics reveals one of its major weaknesses, namely that low-level processes are taken for granted. Much work on *learning robots* has therefore concentrated on learning the details of action execution and its effects, and on learning about the semantic relation between symbolic representations and reality which they represent[1]. Most of this work has been concerned with a single agent. There are some exceptions, however. For example, Alberto Segre's ARMS system learns how to plan from observations of plan executions of a teacher agent. Furthermore, John Laird's *Robo-Soar,* apart from being an application of SOAR to a real world and object manipulation task, permits accepting advice from another agent. In the *World Modellers Project* the goal was to experiment with learning from observations of another agent. The other agent appears solely in the role of a teacher and hence communication between the agents is of a somewhat special kind. Although the investigations into learning robots on a number of important issues, the questions related to communication, cooperation and goal definition have, in general, been left aside. This does not mean that this work is not of potential interest.

Mitchell (1990) describes an autonomous agent (THEO) as having three learning goals: becoming more perceptive, more correct and more reactive. At the moment there is no consensus as to how to map the learning goals to learning methods. Future work could provide some of the answers not only to the questions we have mentioned, but also to the following related issues. When is it better to reason and when to act? When should the system initiate learning and how long should it continue?

[1] The reader could consult (van de Velde, 1991) This collection of papers describes some recent advances in the area of autonomous agent learning.

3. Multi-Agent Learning Systems

Multi-agent systems which include one or more learning agents share some of the concerns of autonomous agent learning. For example, the issue of when to reason, or when to act is even more pertinent in this context. There are important distinctions between the two approaches. Multi-agent learning offers radically different solution to some of the problems in learning. A robot can become more correct (or more reactive) not only by learning from experience, but by communicating with other agents (artificial or human agents). No wonder that the attention of several researchers working in this area has turned to various architectural issues, all of which have something to do with communication. The design should determine *when*, *how* and *with what purpose* should the individual agents communicate. Various systems differ in how they approach these questions. Basically, the learning agents can communicate:

- before the learning / problem solving phase,
- during the problem solving phase,
- before the problem solving phase, but after the individual learning phase,
- during the individual learning phase.

Expressed differently, the agents can be involved in *distributed data gathering, distributed problem solving* or *distributed learning*. Of course various hybrid solutions may exist too.

Distributed Data Gathering + Individual Learning and Problem Solving

Let us finally consider one rather trivial method that enables a number of agents to work on a learning task. All agents are involved in collecting data, but only one system is involved in learning. That is, all the data is transferred to the learning agent that incrementally updates its theory. As the purpose of this paper is not to discuss incremental learning methods, but rather systems with more complex interactions between agents, we shall let the interested reader consult appropriate literature (see e.g. Schlimmer and Fisher, 1988; Utgoff, 1988; Janikow, 1989).

Individual Learning + Distributed Problem Solving

The system described in (Gams, 1989) exploits *redundant knowledge,* and is involved in *distributed problem solving*. It admits several agents with a learning capability, but these do not really communicate while learning is in progress. Different knowledge bases are taken into account when problems are being solved. As has been shown by Gams, this mode achieves a superior performance when compared to a system containing just one knowledge base.

An important issue in this work is how to combine the opinions of different agents. Generally certain confidence factor is associated with each decision and then some method is used to generate the final decision on the basis of the individual decisions.

We notice that distributed solutions need not necessarily involve weighing opinions of different agents. If agent A_i is capable of dealing with a subset of given problems, and if this agent can be considered "sufficiently reliable", we do not need to worry about redundancy at all. The answer of one agent A_i is sufficient. As in Shannon and Weaver's (1964) information theory, the amount of redundancy that is necessary seems to be related to the level of noise present in the data, and the level of uncertainty introduced in its processing. This argument has been put forward by Gams et al. (1990) and is supported by experimental results.

Individual Learning + Knowledge Integration + Individual Problem Solving

The system described by Brazdil and Torgo (1990) attempts to integrate the knowledge acquired by individual agents. The integrated theory is then used by one of the agents to resolve problems.

The system works in three phases. In the first phase the agents go through individual learning. There are no interactions between the agents then. This phase is followed by knowledge integration. This process involves all agents in principle. Knowledge integration can be regarded as a special form of distributed (re-)learning. This process involves characterization of individual theories (or rules) on the basis of experimental tests. These provide the system with estimates of quality or utility of individual theories (rules). This method could be compared with the one used by Gams et al. mentioned earlier employing confidence factors. The quality estimates determine which theories (rules) should be included in the integrated theory.

Experimental results have shown that, in general, the integrated theory had a significantly better performance than the individual theories. We believe that this is due to the fact that redundant knowledge is properly exploited by this system. The knowledge integration method can be seen as a kind of "symbolic filter" for noisy knowledge (imperfect theories and noisy test data).

This approach differs from the one described earlier in several aspects. First, the system can resort to individual problem solving mode. Problems can be directed to the agent that has assembled the integrated theory (although this theory could be given to other agents too).

Problem solving is thus simpler and hence the whole system is more "reactive" if we use the term from autonomous agent learning. It is not necessary to consult the whole community of agents before giving an answer. It is interesting to ask question why this should be so.

As we have mentioned earlier, different agents are called upon many times, but this is done at knowledge integration time. The result of knowledge integration is stored for future use. Consequently one need not consult different agents later. The system of Gams does not attempt to construct such a theory, and so it is necessary to solicit opinions of other agents at problem solving time.

There are arguments for and against each approach. The system described by Gams retains structured representation of knowledge. As individual agents update their knowledge, this immediately bears some effects on the opinion of the group. This is not true of the integrated theory. If one of the individual theories has been altered, the integrated theory may need to be revised. In a certain sense, the first approach has similar advantages as *interpreting*, while the second one has the advantages of *compiling*.

Sometimes it may be difficult or outright impossible to construct an integrated theory. Difficulties can arise particularly when the agents use different (and possibly incompatible) ways of representing knowledge.

Integrated theory represents a more compact representation of knowledge than the structured representation discussed earlier. Compact representations have obvious advantages. Simple theories are easier to communicate to other agents (including humans) than complex ones. They can also serve as a useful starting point in further learning.

Alternative theories are no doubt useful both in science and politics. Alternative theories often find their adepts, and it would be wrong to try to come up with one integrated theory that would explain everything. However, people would generally agree that there is a limit as to how many theories should be taken into account. Some theories may be just minor variants of others. In our view methods are needed that would determine whether some particular theory is worth keeping around as a useful alternative.

Distributed Learning + Individual Problem Solving

Sian's system (1990a, 1991b) is involved in both individual and distributed learning. Each system learns individually, but if certain conditions arise interaction is initiated with other agents. The interaction is established via an *interaction board*, which plays a similar as in blackboard architecture systems. Here the agents can, for example, propose a hypothesis to the the interaction board.

Communication between the learning agents is whenever one of the agents has obtained a hypothesis and has sufficient confidence in it. This is considered as a good candidate to put to test. Opinions of the other agents are solicited with the objective of establishing a consensus. The rules can remain as they are, or they can be modified, or they can be withdrawn. The rules that have been agreed upon represent a *consensus* of the group and appear in the "integrated theory".

This work differs form the other two presented earlier in various aspects. First, the author has elaborated an interface through which the individual agents communicate. Introduction and retraction of hypotheses to/from the interaction board is achieved using the operators

PROPOSE, ASSERT, WITHDRAW, ACCEPT

Evaluation of hypotheses is accomplished using the operators

CONFIRM, DISAGREE, MODIFY, NOOPINION,

while AGREED modifies a state. Each hypothesis is characterized by a NET-VALUE calculated on the basis of the opinions of different agents (CONFIRM, DISAGREE, MODIFY, NOOPINION) and the confidence values associated with each operator.

We notice that all three systems discussed in this section (i.e. Gams's, Brazdil & Torgo's and Sian's) use some particular method for assessing the usefulness of a given rule on the basis of evidence presented by different agents. Further work could be done to present a more detailed comparative study.

As we have mentioned earlier Sian's system differs from Brazdil and Torgo's in one important aspect. The agents are allowed to interact in the learning phase. This seems to make sense, particularly if we are interested to save some agents' effort associated with learning. The earlier a potentially good hypothesis is put to test and possibly accepted, the better.

When considering testing in a multi-agent environment, it is necessary to distinguish between centralized testing (done by one agent) and distributed testing. Testing against all data available does not necessarily imply a centralized solution. A particular hypothesis can be sent to different agents. Each can then update the information received. A global view can be thus built up by passing a hypothesis from one agent to another.

In Sian's system each agent tests the proposed rule against his *own data.* A global view of each rule is then built up on the basis of a number of local views. A question arises whether this built-up view is equivalent to the global view that could be obtained by centralized testing. Brazdil and Torgo's system seems to satisfy this criterion. Each agent could update the qualitative and quantitative characterixation of the given rule and then pass this information to the next agent. This information is the same as the one generated during centralized testing.

Further work could be done here. A study of cost-effectiveness of the two methods could be made, taking into account:

- the effort of transferring the local views / instances to one agent,
- the effort of evaluating a given hypothesis (using instances / local views),
- net increase of confidence for some particular method.

4. Some Aspects of Communication between Agents

As has been suggested in the previous sections, communication plays rather an important role in multi-agent learning systems. It may supply the agent with valuable information and thus avoid "re-discovering the wheel". Communication need not, however, bring about benefits. It is thus important to study this topic in its own right. Although this topic exceeds the objective of this paper, we would like to make several observations here.

It is important to distinguish between the issues related to *form of the language* used between agents and the actual *statements* in that language. Here we make a similar distinction as when talking about natural language. There is a difference between problems related to structure of English and particular piece of text.

The issues related to the language itself could be viewed as issues of interfaces between agents. It is necessary to decide what kind of statements the agents should be able to generate and comprehend. For example, one could decide that the operator PROPOSE(H,C) should have a certain meaning. In Sian's system this operator adds hypothesis H (and the associated confidence C) to the interaction board.

Obviously, the design of interfaces is closely related to the design of the architecture of the whole system. The operator PROPOSE plays a specific role in the system for which it was designed. A question arises whether some set of basic communication primitives could be found that would be generally useful in multi-agent learning. This would have the advantage that it would make it easier to compare different approaches. Of course, one could always add extra primitives, or define other constructs in terms of the existing core primitives, if this was required in some specific system.

The second kind of issues are related to the problems of interpretation and meaning of agent's statements. As Shaw and Gaines (1989) have pointed out, same term can have different meaning for different agents. This situation is called a *conflict*. Different terms may, however, have similar meaning. This situation is called a *correspondence*.

Work of Brazdil and Muggleton (1991) is concerned with the problem of resolving certain language differences between agents. The agents are not only presented with different situations from which they can learn, but also, employ a somewhat different terms in their description of the (simulated) world. For example, one agent uses the predicate *father(..)* while the other *parent(..)*. If we use Shaw and Gaines's terminology, there is a problem of correspondence. Brazdil and Muggleton show how these language differences can be overcome. It is shown that standard machine learning techniques can be used to acquire the meaning of undefined concepts.

There are interesting relationships between inductive learning and communication. There interplay mentioned here is of a different kind than the one discussed in Section 3. There we have discussed different ways communication can supplement learning. Here we are concerned with the possibility of resolving certain problems of communication using learning.

5. Role of Learning in a Community of Agents

Utility of Learning

Learning in distributed systems opens new horizons. It makes us consider issues that have not been looked at earlier in machine learning. For example, it forces us to consider the question why a particular agent (in a community of agents) should want to learn? Designing agents that would learn about anything in the world goes against the basic philosophy of distributed AI.

We believe it is thus necessary to reason about the *utility of learning*. We notice that in most general architectures of intelligence (SOAR, THEO, PRODIGY, ICARUS) this issue has not really been paid attention to. This may be the reason why some systems are ill-behaved (the more they learn, the slowed they perform). We believe that addressing this point in the context of DAI will make it easier to find the appropriate answer(s).

Community Goals and Agent Goals

An important capacity of an agent in a multi-agent world is the ability to define one's own goals. The agent's goals are often affected by (and in some cases determined by) the goals of other agents.

The process of *goal definition* seems to subsume goal selection. Mitchell (1990) has defined *perception* as a process linking the state of the world to the appropriate goals to attend to. This process involves selecting the most pertinent goal and trying to achieve it in preference to others.

So far not much work has been done in the area of goal definition. Most work done has concentrated on goals of two agents only. Baker's system KANT (1991), for example, incorporates reasoning mechanisms for determining which set of goals are to be *negotiated* in a tutorial interaction.

When considering the relationships between individual goals and community goals two issues arise. First, how the satisfaction of individual goals affects the satisfaction of community goals. Then, how the satisfaction of community goals leads to the satisfaction of individual goals.

It is also possible to envisage that agents could *learn which goals to pursue* in order to achieve some overall goals. Perhaps the agents would follow a scheme of gradual differentiation that is common in human society. The agents start with similar goals, but differences in local conditions and agent-specific skills gradually differentiate the agents' goals so as to function better in a community. Ideally this process lets the community evolve from a fairly uniform group to a differentiated highly competent society.

Learning Tasks in a Community

As we have mentioned earlier, Mitchell (1990) ascribes three learning goals to an agent: becoming more perceptive, more correct, and more reactive. These are the goals that an external observer might ascribe to an agent when observing its behaviour over time.

If an observer were to observe a multi-agent system, which learning goals could he ascribe to individual agents? We believe that the list of learning goals mentioned by Mitchell could be extended to include at least one additional requirement. We could require that the agent should become increasingly *more integrated,* that is, play its proper role in the community of agents.

This involves being called upon by other agents, recognizing when to delegate a problem to others and exploiting these opportunities.

Acknowledgments

The work of P. Brazdil and L. Torgo discussed in this paper was supported by Esprit 2 Project Ecoles (3059). The authors wish to thank the Commission of European Communities for their support.

References

Baker M. (1991): "A Model for Negotiation in Intelligent Tutoring Dialogs", in *New Directions for ITS*, E.Costa (ed.), to be published by Springer Verlag.

Bond A. and Gasser. L. (eds.) (1988): *Readings in DAI*, Morgan Kaufmann Publishers, San Mateo.

Brazdil P. and Torgo L. (1990): "Knowledge Acquisition via Knowledge Integration", in *Current Trends in Artificial Intelligence, B. Wielinga et al. (eds.)*, IOS Press, Amsterdam, 1990.

Brazdil P. and Muggleton S. (1991): "Learning to Relate Terms in Multiple Agent Environment", in this volume.

Buntine W. (1889): "Learning Classification Rules Using Bayes", in *Proceedings of 6th International Workshop on Machine Learning*, Ithaca, New York.

Demazeau Y. and J.-P. Mueller (eds.) (1991): *Proceedings of the 2nd European Workshop on Modelling Autonomous Agents and Multi-Agent Worlds (MAAMAW 90)*, Saint-Quentin-en-Yvelines, France, August 1990, Elsevier Science Publishers.

Durfee E., Lesser V.R. and Corkill D.D. (1989): "Cooperative Distributed Problem Solving", in *The Handbook of Artificial Intelligence, Volume IV*, Barr A., Cohen P.R. and Feigenbaum E.A. (eds.), Addison Wesley, 1989.

Gams M. (1989): "The Measurement Highlight the Importance of Redundant Knowledge", in *Proceedings of 4th European Working Session on Machine Learning (EWSL-89)*, K. Morik (ed.), pp. 71-80, Pitman - Morgan Kaufmann.

Gams M., Bohanec M. and Cestnik B. (1990): A Schema for Using Multiple Knowledge, Working Paper, Jozef Stefan Institute, Ljubljana, Yugoslavia.

Huhns M. (ed.) (1987): *Distributed AI Vol.1*, Pitman and Morgan Kaufmann, 1987.

Janikow C.Z. (1989): "The AQ16 Inductive Learning Program: Some Experimental Results with AQ16 and Other Symbolic and Nonsymbolic Programs", Rep. of AI Center, George Mason University.

Mitchell T. (1990): "Becoming Increasingly Reactive", in *Proceedings of 8th National Conference on Artificial Intelligence (AAAI-90)*, Morgan Kaufmann, pp. 1051-1058.

Oliveira E., Camacho R. and Ramos C. (1990): A Multi-Agent Environment in Robotics, Working Paper, Fac. of Engineering, Univ. of Porto, (also LIACC, Univ. of Porto), Portugal.

Schlimmer J.C. and Fisher D. (1986): "A Case Study of Incremental Concept Induction", in *Proceedings of the Fifth National Conference on Artificial Intelligence*, pp. 496-501, Morgan Kaufmann.

Sian S. (1991a): "Adaptation Based on Cooperative Learning in Multi-Agent Systems", in *Proceedings of the 2nd European Workshop on Modelling Autonomous Agents and Multi-Agent Worlds (MAAMAW 90)*, Demazeau Y. and J.-P. Mueller (eds.), Saint-Quentin-en-Yvelines, France, August 1990, Elsevier Science Publishers.

Sian S. (1991b): Extending Learning to Multiple Agents: Issues and a Model for Multi-Agent Machine Learning (MA-ML), in this volume.

Shaw M. and Gaines B. (1989): Knowledge Acquisition: Some Foundations, manual Methods and Future Trends, in *Proceedings of Third European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, J. Boose, B. Gaines and J.G. Ganascia (eds.), Paris, July 1989.

Utgoff P.E. (1988): "ID5: An Incremental ID3", in *Proc. of 5th International Workshop on Machine Learning*, J. Laird (ed.), Ann Harbour, Morgan Kaufmann Inc.

van de Velde W.(ed.) (1990): *Towards Learning Robots*, North Holland, Amsterdam.