KNOWLEDGE ORIENTED LEARNING

Paul D. Scott and Robert C. Vogt

Department of Computer and Communication Sciences, University of Michigan

ABSTRACT

Two alternative learning strategies are defined and discussed: Task-oriented learning in which a system tries to improve its performance at a specified task and knowledge-oriented learning in which a system builds an organised representation of experience. It is argued that the relationship between these two strategies is analagous to the relationship between technology and science. Hence it is concluded that the development of knowledge-oriented learning systems, which has been largely overlooked by the AI community, is a worthwhile research goal.

We then proceed to describe how such a system may be constructed by building a machine which continually tries to reduce its own uncertainty regarding the outcome of its actions. An implementation which learns to perform a multiple concept learning task in the presence of noisy data is briefly described.

1: INTRODUCTION

At present the successes achieved by various knowledge-based AI systems can only be attained through a considerable effort by human experts who supply the necessary knowledge in a suitably organized form. The notion of a knowledge-based system endowed with the ability to acquire its own knowledge base is thus becoming increasingly attractive.

Much of the research on machine learning [1] appears to be based on the same basic assumption regarding the nature of learning. This assumption has been succinctly expressed in Simon's [6] definition of learning as any process by which a system improves its performance. Note that this defines what learning is for, and thus views learning as any means for achieving a particular end, rather than attempting to specify what learning actually is.

*We thank the NSF and Horace H. Rackham School of Graduate Studies for financial support.

Alternative definitions are possible. One could for example define learning as the construction of an organized representation of experience. This definition is complementary to Simon's since it specifies the means but makes no reference to the end. The two definitions provide alternative viewpoints from which the same system can be regarded.

These alternative viewpoints suggest different approaches to constructing a learning machine. Simon's definition suggests a system which repeatedly attempts to perform some task and, on the basis of the results of each attempt, makes some changes in its internal organisation. The effect of these changes should be a steady improvement in the systems's performance at that task. Most machine learning systems follow this paradigm which we shall term 'Task-Oriented Learning'.

Defining learning as the construction of an organized representation of experience suggests a different kind of learning machine. This would be a system which repeatedly interacted with its environment and built an internal model which summarised its experiences. In such a system there would be no attempt to improve performance at any particular task. We shall term this 'Knowledge-Oriented Learning'.

Is it possible to construct a knowledge-oriented learning system? Could it learn anything useful? Would it ever have advantages over a task-oriented learning system? Under what circumstances? The remainder of this paper attempts to answer these questions.

2: AN ANALOGY

We begin with an analogy. Science is a process by which people have collectively developed a logical, integrated and self-consistent model of events which occur in the physical world. It has two complementary aspects. One is technology: those investigations designed to achieve a particular practical goal. The other aspect is pure science: those investigations intended only to extend our understanding of the way the observable world works. Even a slight acquaintance with the development of scientific thought reveals that these two aspects are inseparable. Each is essentially dependent on the other.

Technology is society operating a task-oriented learning strategy. Similarly pure science is society operating a knowledge-oriented learning strategy. Since the two are complementary processes in society we should seriously examine the possibility that the same may be true for machine learning systems. We shall therefore pursue the analogy in order to clarify the role of knowledge-oriented learning.

Suppose that in 1800 some government had asked its scientific academy to invent a device for transmitting pictures rapidly over large distances. This hypothetical task-oriented learning assignment illuminates several important points.

First it is absurd since the project was not technically feasible in 1800 and would have been regarded as impossible. The pure scientific endeavours of Faraday, Maxwell, Hertz and Thompson were necessary precursors of television. Only when knowledge-oriented learning had created a rich enough model of the physical world could task-oriented learning be used to achieve the practical end.

Second, it is doubtful if those eminent physicists would have done the research they did if their goal had been a remote viewing system. It is even more doubtful that a government concerned to acquire such a system would have continued to fund such theoretical work. The results obtained by knowledge-directed learning do not necessarily yield results which are immediately or obviously useful.

Third, the knowledge of electromagnetic phenomena obtained in the nineteenth century has not only served as the basis for one particular invention but is fundamental to much of today's technology. The long term benefits of knowledge-oriented learning may have a much more fundamental utility than the results of task-oriented learning.

This analogy is open to several objections. First, fundamental advances in pure science were achieved by some of the world's most intelligent people. AI has not reached the stage where we can endow our creations with such abilities. This objection is countered by observing that all humans and some animals make extensive use of knowledge-oriented learning during their early lives in order to acquire motor, perceptual, cognitive and social skills. The same strategy which is dignified as 'science' when it occurs in a laboratory is dismissed as 'play* when observed in a nursery.

A second objection would argue that while knowledge-oriented learning is needed for fundamental advances it has a negligible role in most of our learning activities. Like the previous objection this may be refuted by the extensive use of play as a learning strategy. Indeed this suggests that knowledge-oriented learning yields benefits which are not achievable by other means. (For further discussion of play viewed as knowledge-oriented learning see [3] and [14]) .

A third objection would be that we want learning machines to learn something we regard as useful. Any system which acquires knowledge without an externally specified goal will either build a representation of something irrelevant to the user's needs or, more likely, accumulate a morass of unrelated observations and spurious correlations. We propose to refute this objection by describing the construction of a system which uses knowledge-oriented learning to build a useful representation of its environment.

3: SEARCHING FOR CERTAINTY

As noted earlier, science attempts to construct a logical, integrated and self-consistent world model. This is subject to continual correction so long as it remains either incorrect or incomplete. If it is incorrect it leads to erroneous predictions. If it is incomplete there are situations whose outcome cannot be predicted. Thus the model needs modifying whenever something unexpected happens.

This principle may be used as the basis for building a machine which will construct a representation of its environment. Such a machine would interact with its world in an organised and purposeful fashion. Its underlying goal would be to avoid being surprised, not by avoiding novelty but by constructing a model which enabled it to predict the outcome of all its actions.

This is a special kind of goal. In seeking an ordinary goal a system is attempting to change its environment. To do this it uses a representation of that environment. In contrast, in constructing a model to eliminate surprise a system is trying to change its representation of the environment. To do this it would use knowledge about that representation. For example it must have some representation of its own ignorance: the region of uncertainty on the borders of its knowledge. Such a representation would be a form of metaknowledge.

This metaknowledge need not be sophisticated to form the basis of a successful knowledge-oriented learning system. It could take the simple form of associating a measure of certainty with each item in the knowledge base. The system could then operate as follows:-

1. Select least certain item in knowledge base.

2. Perform some action whose outcome will substantiate or discredit that item. (ie. do an experiment)

3. Change knowledge base, including certainty measures, in accord with outcome of experiment.

4. Goto 1

Such a system would continually reduce the uncertainty in its original representation. However the changes introduced in step 3 may well give rise to further areas of uncertainty which will themselves be subsequently reduced.

Note that the system must start with an initial representation. This could be either very simple, as in the example following, or very elaborate, in which case the learning system would proceed to refine it.

4: PAN : AN IMPLEMENTATION

We have written a program called PAN to demonstrate that a knowledge-oriented learning system can be constructed using these ideas. Space limitations restrict us to sketching some of the main features of the system but a detailed account will be found in [5]. PAN is a simulated robot operating in a simple environment called PANWORLD which consists of a collection of objects. PAN can distinguish these objects on the basis of various attributes such as color, size, texture and shape. In addition PAN can detect when an object moves.

PAN can perform five different operations on individual objects: Push, Pick-Up, Kick, Zap and Woggle. The outcome of an operation is determined by a set of rules called PANPHYSICS which are chosen by the experimenter. These rules may be either deterministic or stochastic. A typical deterministic rule is:-

> Only objects which are large and rough move if woggled.

A typical stochastic rule is:-

> Probability that smooth object moves if pushed is 0.8.
> Probability that non-smooth object moves if pushed is 0.3.

Stochastic rules would create serious problems for many concept learning systems but they are readily discovered by PAN.

PAN attempts to construct a representation which will enable it to predict the outcome of applying any of its operations to any object. Such knowledge is of obvious utility in problem-solving. It does this by developing a hierarchy of classes of objects. Initially this hierarchy comprises the single class 'Things' of which every object is an instance. Subsequent classes are subclasses of this class and membership comprises those objects possessing a distinguishing set of attribute values.

Attached to each class is a set of operations. These are those operations which PAN believes will result in any object in that class moving. This representation was deliberately modelled on object-based languages such as Smalltalk. Associated with each action in each class is an estimate of the probability that an instance of that class will move if that operation is applied to it. Initially all operations are attached to the class 'Things' with associated probability 0.5.

The measure of uncertainty used to drive PAN is derived from the probability estimate. We used Shannon's information function [6] in anticipation of extending the system to experiments with many possible outcomes but other functions could be used to produce similar results.

The operation of PAN follows the scheme discussed above. First the least certain item in the knowledge base is chosen. This is the operation which is attached to a class with highest uncertainty. An experiment is then performed. The operation is applied to a randomly chosen instance of that class. The object will either move or remain stationary.

If it moves the current representation appears to be correct so the only change made is an increment in the corresponding probability estimate.

If the object does not move the system has generated an incorrect prediction. The corresponding probability estimate is decremented. The class hierarchy is also modified in one of two ways. If the resulting probability is lower than the probability with which the operation is attached to the superclass of the class then the operation is detached from the class. Classes without subclasses or attached operations are deleted.

If however the resulting probability is still higher than that in the superclass PAN concludes that some but not all of the current class will move when the operation is applied. It therefore attempts to construct a subclass of the current class, all of whose members move when the action is applied. It applies the operation to instances of the current class until it finds one that moves. An attribute value of this object is added to the attribute values which distinguish members of the current class to form the distinguishing attribute values for the new subclass. The action is then attached to the new subclass.

The system then begins a new cycle. This process continues until all the uncertainty values in the system drop below a threshold value which typically corresponds to probability estimates of 0.05 and 0.95. At this point PAN has a great deal of confidence about the likely outcome of any experiment. Since it does not seek absolute certainty and since its only memory of previous experiments is in the form of probability estimates it is not sensitive to low levels of indeterminism in PANPHYSICS. Higher levels of indeterminism prevent the system reaching its certainty thresholds but the system still builds a representation which correctly reflects the intrinsic uncertainty of its environment.

5: CONCLUSION

PAN is a simple system which demonstrates that a knowledge-oriented learning strategy, based on knowledge of its own uncertainty, works well on a multiple concept learning task in the presence of noisy data.

It is interesting to compare it with another much more sophisticated learning program: AM (see [2]). AM starts with some knowledge of finite sets and proceeds to discover much of number theory. It is a knowledge-oriented learning system since it is not trying to improve its performance. It simply explores those combinations which appear 'interesting'. AM includes 59 heuristics to assess 'interestingness' and these play the same role as the uncertainty measure in PAN.

There is clearly considerable scope for the development of both other pure knowledge-oriented systems and hybrid systems which exploit both learning strategies.

6: REFERENCES

[1] Cohen,P.R. & Feigenbaum,E.A., "The Handbook of Artificial Intelligence : Volume 3". Kaufmann, Los Altos, Calif., 1982.

[2] Lenat,D., "AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search", Doctoral dissertation, Stanford University, July 1976.

[3] Scott,P.D., "Play Considered as a Strategy for Knowledge Acquisition", Proc. 4*th* Ann. Conf. Cognitive Science Society, Ann Arbor, USA, August, 1982.

[4] Scott,P.O. 6 Vogt.R.C, "On Play and Learning", In preparation.

[5] Scott,P.D. 6 Vogt.R.C, "An Approach to Knowledge-Oriented Learning", In preparation.

[6] Shannon,C.E. and W.Weaver, "The Mathematical Theory of Communication", University of Illinois Press, Urbana, Chicago, London, 1949.

[73 Simon,H.A., "Why should machines learn?", In "Machine Learning", Eds R.S.Michalski, T.M.Mitchell & J.Carbonell, Kaufmann, Los Altos, Calif., 1983.