Adaptive Interactions in Societies of Agents

Brian R. Gaines
Knowledge Science Institute
University of Calgary
Alberta, Canada T2N 1N4
gaines@cpsc.ucalgary.ca

Abstract

In a society of agents the learning processes of an individual agent can become critically dependent on knowledge processes in the society as a whole. It is useful to adopt a *collective stance* that views the overall society as a larger adaptive agent that is recursively divided into adaptive sub-agents. The resources for an agent include other agents, and some of an agent's processes will be devoted to modeling other agents' capabilities and others to developing those capabilities. This article presents a model for adaptive interactions in societies of adaptive agents in which *knowledge* arises as a state variable imputed by one agent to another to account for its capabilities, and task allocation between agents results in *functional differentiation* in an initially homogeneous society. It is shown that a simple *training* strategy of keeping an agent's performance constant by allocating tasks of increasing difficulty as an agent adapts optimizes the rate of learning and linearizes the otherwise sigmoidal learning curves. It is suggested that this is the basis for the human preference for a *flow* condition in which the challenge of a task is managed to remain between the extremes of boredom and anxiety.

1 Introduction

Early research on *learning machines* (Gaines and Andreae, 1966) saw these machines as computational modules that could be programmed indirectly through experience, and focused on the manipulation of the modules through *coding*, *priming* and *training* (Gaines, 1968).

Coding is concerned with appropriate input and output interfaces which are known to be critical to the learning capabilities of both machines and people—minor changes in information encoding can change a task from one which is very easy to learn to one which is virtually impossible.

Priming is concerned with the transfer of knowledge not through learning from experience but through other mechanisms such as mimicry, analogy and language (Gaines, 1969). It was shown that mechanisms for the linguistic transfer of control strategies to perceptron controllers could decrease their learning time in a way similar to that for human controllers given the same instructions (Gaines, 1972a), and make the difference between a task being very easy or virtually impossible to learn. These result led to experiments with the use of linguistic *fuzzy hedges* (Zadeh, 1973) for priming learning machines with the surprising result that priming alone was sufficient to induce excellent performance in some situations (Mamdani and Assilian, 1975), triggering the development of fuzzy control (Sugeno, 1985).

Training is concerned with the sequence of learning environments presented to the learning system in order to maximize its rate of learning. It was shown that for both learning machines and people a 'feedback trainer' that adjusted the difficulty of a task to keep performance constant greatly increased the speed of learning, and that tasks which were virtually impossible to learn under conditions of fixed difficulty could be learned rapidly with feedback training (Gaines, 1968; Gaines, 1972b).

The effects of coding, priming and training are basic phenomena in any learning system. Now that the aspirations of the 60s are beginning to be filled by the intelligent adaptive agents of the 90s, it is timely to revisit some of these phenomena. In particular, research in distributed artificial intelligence (Bond and Gasser, 1988), artificial life (Langton, 1995), and cultural evolution (Boyd and Richerson, 1985), raises issues of how a society of agents interacts to provide one another with mutual training environments.

This article describes a collective stance model of learning phenomena in agent communities that accounts for the development of functional differentiation in uniform populations of adaptive agents. In particular, it develops a quantitative model of the way in which the sigmoidal learning curves of agents learning in static environments can be linearized through feedback training, and uses this model to account for a range of phenomena in communities of intelligent adaptive agents.

2 Taking a Collective Stance

A useful perspective from which to examine learning phenomena in agent societies is a *collective stance* (Gaines, 1994) in which the society is viewed as a single adaptive agent recursively partitioned in space and time into sub-systems that are similar to the whole. In human terms, these parts include societies, organizations, groups, individuals, roles, and neurological functions (Gaines, 1987). Notions of expertise arise because the society adapts as a whole through the adaption of its individual agents. The phenomena of expertise correspond to those leading to distribution of tasks and functional differentiation of the individual agents.

The mechanism for functional differentiation is one of positive feedback from agents allocating resources for action to other agents on the basis of those latter agents past performance of similar activities (Gaines, 1988). Distribution and differentiation follow if performance is rewarded, and low performers of tasks, being excluded by the feedback mechanism from opportunities for performance of those tasks, seek out alternative tasks where there is less competition.

The knowledge-level phenomena of expertise, such as meaning and its representation in language and overt knowledge, arise as byproducts of the communication, coordination and modeling processes associated with the basic exchange-theoretic behavioral model. The collective stance model can be used to account for existing analyses of human action and knowledge in biology, psychology, sociology and philosophy (Gaines, 1994).

Simple simulation experiments of a competitive environment for two agents can illustrate the formation of expertise through positive feedback (Gaines, 1988). For example, let the rules of a basic phenomenological simulation be that: each problem requires certain knowledge; if the agent does not have the knowledge necessary it guesses with a probability of success, learning if it succeeds; the society chooses the expert for a problem with equal probability initially, gradually biasing the choice according to success or failure; there is no communication of knowledge between experts. Figure 1 graph A plots the probability that one agent will be always preferred and shows that this rapidly approaches 1.0—a best 'expert' is determined. Graph A shows the expected knowledge of the better expert and graph B that of the rival—one goes to 100% rapidly and the other is asymptotic to about 36%—there is objective evidence of the superior ability of the chosen expert. Which of the two agents becomes the 'best expert' is, of course, completely chance.

The simulation shown is not Monte Carlo but based on the calculation of the exact probability distributions involved. It can easily be adjusted to take into account differences between the experts: that one starts with greater knowledge; that one learns faster; that one is favored initially (the prima facie credibility or 'well-dressed consultant' phenomenon). Similar simulations have been made of different positive feedback mechanisms; for example, if both experts are given the same problem but the problem difficulty is adjusted upwards if either gets it right—the situation of keeping up in the scientific 'rat race.' Effects have been introduced of the loss of knowledge through inadequate opportunities for its application, the growth of scientific knowledge so that there is always more to be acquired, and differential access to priming through cultural knowledge transfer processes such as education. All the simulations bear out the expected qualitative result, that a range of different positive feedback mechanisms in an agent society are adequate to account for differential expertise in a sample with initially equal knowledge and capabilities.

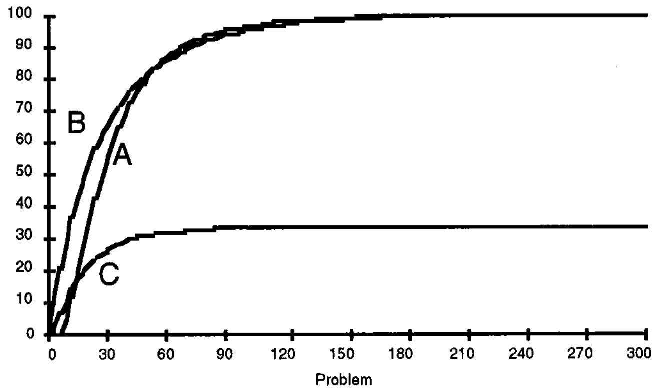

A: Percentage probability that problem will be given to only one of the two 'experts'
B: Expected knowledge of superior 'expert'
C: Expected knowledge of inferior 'expert'

Figure 1 Simulation of the effects of positive feedback on the formation of expertise

There is evidence that these processes occur in human communities. Sociologists have noted positive feedback processes in the dynamics of the scientific community (Hagstrom 1965). Merton (1973) coined the term the "Matthew effect" for those features of the reward system in research that were biased towards allocating greater credit for the same discovery to those with an already established reputation. In medicine, a key learning resource is access to medical problems, and the 'owner' of such a problem has a keen personal interest in only allowing someone of very good reputation to handle it. The system, including considerations of legal liability, funnels problems to those who are regarded as experts. It is, however, precisely these problems from which new knowledge is generated. Similar considerations apply to the award of scholarships, invitations to scientific congresses, and so on (Blume 1974). They also apply not only to individuals but also to social units such as a company subject to government procurement procedures that are heavily biased to contractors with 'prior experience' and with whom the government agency has 'prior experience.'

3 Performance-Based Feedback Linearizes the Learning Curve

The effect upon individual agents of the social feedback processes described in the previous section is to regulate the learning experiences available to the agent in such a way that agents with historically better performance get the tasks judged to be more difficult. In a world with a wide range of tasks of varying difficulty and with agents having limited capacities and lives, the overall effect is that each agents performance is kept roughly constant as it is given tasks of increasing difficulty commensurate with its learning as indicated by its performance on previous tasks. The agent is learning to cope with tasks of increasing difficulty so that its skills are increasing, but being constantly presented with tasks just beyond its capabilities so that its performance is not.

It is interesting to attempt to develop a qualitative model of the learning phenomena involved using as weak assumptions about the adaptive agent as possible so that the model is widely applicable. Consider a universe of *tasks* and a universe of *knowledge* such that an agent can perform a task if it has some, not necessarily unique, collection of knowledge. Assume (it will be argued later without loss of generality) that items of knowledge form a set-theoretic structure, and that the *difficulty* of a task can be estimated in terms of the cardinality of a set of knowledge that allows an agent to perform it.

Assume that when an agent is given a task for which its knowledge is inadequate that the probability of it guessing each missing item of knowledge is q and that if it guesses all the missing items it performs the task and learns the knowledge, but otherwise learns nothing.

Assume that the training system can select a task at a given level of difficulty but does not know either the knowledge required to perform it or the state of the agent's knowledge.

Assume that the probability that the agent knows a randomly selected item of knowledge is p, and that the probability that it learns an unknown item during a task necessitating it is q.

Then the probability that an agent will perform a task of difficulty d is:

$$P(n) = \sum_{0}^{d} \binom{d}{i} p^{d-i} (1-p)^{i} q^{i} = \left(p + (1-p)q\right)^{d} \tag{1}$$

and the expected rate of learning is:

$$L(n) = \sum_{0}^{d} \binom{d}{i} i\, p^{d-i} (1-p)^{i} q^{i} = d(1-p)q\left(p + (1-p)q\right)^{d-1} \tag{2}$$

If one selects a task difficulty that optimizes the rate of learning by maximizing L(n) with respect to n then:

$$d_{opt} = -1/\ln(p + (1-p)q) \tag{3}$$

and the expected performance to achieve this is:

$$P_{opt} = 1/e \approx .37 \tag{4}$$

and the expected maximum rate of learning is:

$$L_{opt} = -(1-p)q/(e(p + (1-p)q)\ln(p + (1-p)q)) \tag{5}$$

which is such that as agent learns and p approaches 1:

$$\underset{p \to 1}{L_{opt}} = q/e(1-q) \tag{6}$$

The implication of equation (4) is highly significant because it implies that a training system that adjusts the task difficulty to keep performance constant will achieve a linear rate of learning that is the fastest possible.

One can also deduce from equation (4) that the optimum performance involves the learning agent being correct 37% of the time—i.e. an error rate of 63%. However, this result is misleading unless one does a sensitivity analysis. Such an analysis shows that the learning rate drops to half optimum when the error rate decreases to 20% or increases to 80%—i.e. that the optimum learning performance is relatively insensitive to the performance set-point chosen. The reason for this is that there is a lower chance of learning more at lower performances, and a greater chance of learning less at higher performances.

Thus a feedback trainer that adjusts task difficulty to maintain performance constant can induce optimum learning in an adaptive agent, and achieve something close to this even if the performance fluctuates over a 5 to 1 range.

What happens if an adaptive agent is given tasks to perform of constant difficulty. One would expect the learning rate to be sub-optimal initially because the task is to difficult, become optimal for a period as the performance improves to a level where the difficulty is optimal, then decline again as the performance improves to a level where the task is too easy. Figure 2 shows this results of integrating equation (2) for fixed n compared with the results of integrating (5), with q = 0.5 and full knowledge being 25 items. It can be seen that feedback training achieves a speedup in learning by a factor of 2 compared with the best fixed-difficulty training, that the learning curves have the expected sigmoidal shape, and that training on tasks of high difficulty involves very long learning periods.

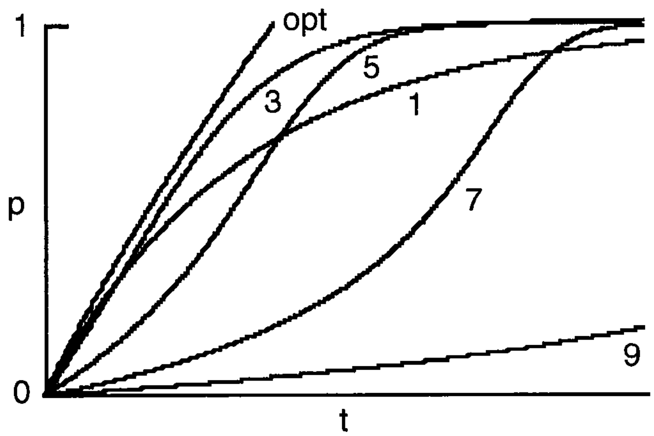

Figure 2 Learning curves for adaptive agent under various training conditions

In practice there may be complications that make the effects shown even more pronounced. For example, the universe of simple tasks may be incomplete in that some knowledge items may only be brought into use in relation to more complex tasks. Such dependencies between knowledge items make feedback training essential.

3.1 Corollary 1: Knowledge as an Imputed State Variable

The generality of the results of Section 3 are dependent on the general applicability of the learning model used. It was noted that it would be argued that the assumption that knowledge is a set of items can be made without loss of generality. The basis for this assertion is that the notion of "knowledge" is itself ill-defined and open to analysis and definition. From a situated cognition perspective, knowledge is some index of adaptivity:

> "The new perspective, often called *situated cognition*, claims that all processes of behaving, including speech, problem-solving, and physical skills, are generated on the spot, not by mechanical application of scripts or rules previously stored in the brain. Knowledge can be represented, but it cannot be exhaustively inventoried by statements of belief or scripts for behaving. Knowledge is a capacity to behave adaptively within an environment; it cannot be reduced to representations of behavior or the environment." (Clancey, 1989)

The representation of the skills of an agent by imputing to it "knowledge" as a *state variable* indexing those skills is a social phenomenon (Gaines, 1989). It is the way in which we model one another for the purposes of task allocation and skill development. It is reasonable to hypothesize that our models are effective in terms of task indexing and training, and to reverse the arguments of Section 3 to propose that our folk models of knowledge as a "substance" which is "possessed" by an agent are founded upon the results of that section—i.e. under this model, optimum training leads to linear increase in knowledge with time—the natural units of knowledge are those of improvements per unit time in adaptive behavior.

3.2 Corollary 2: Flow as an Optimum Learning State

This article has emphasized the social provision of the optimal training conditions for learning. However, what motivates the learner to seek out these conditions? Csikszentmihalyi's (1990) concept of *flow* as the phenomenon underlying the psychology of optimal experience provides a model of the learner dynamics. Hoffman and Novak (1995) summarize the concept as:-

> *"Flow has been described as 'the process of optimal experience' achieved when a sufficiently motivated user perceives a balance between his or her skills and the challenges of the interaction, together with focused attention."*

The likeability of a task correlates with a flow state in which a motivated user undertakes a task whose level of difficulty is at some particular level that suits their individual needs. Too low a level results in boredom and too high a level in anxiety, and the optimal level results in the intense satisfaction with the activity that Csikszentmihalyi terms *flow.*

As the agent learns a flow state can be maintained only by increasing the task difficulty to keep the performance constant. As shown in Section 3, this also maximizes the rate of learning, and this suggests that the flow process may have evolved phylogenetically as a mechanism reinforcing an agent which is maximizing its rate of learning. If the reinforcement center in the brain is stimulated by conditions that optimize learning, then individuals will be attracted to socially created learning environments.

4 Conclusions

In a society of agents the learning processes of an individual agent can become critically dependent on knowledge processes in the society as a whole. It is useful to adopt a *collective stance* that views the overall society as a larger adaptive agent that is recursively divided into adaptive sub-agents. The resources for an agent include other agents, and some of an agent's processes will be devoted to modeling other agents' capabilities and others to developing those capabilities.

This article has described a model for adaptive interactions in societies of adaptive agents in which *knowledge* arises as a state variable imputed by one agent to another to account for its capabilities, and task allocation between agents results in *functional differentiation* in an initially homogeneous society.

It has been shown that a simple *training* strategy of keeping an agent's performance constant by allocating tasks of increasing difficulty as an agent adapts optimizes the rate of learning and linearizes the otherwise sigmoidal learning curves.

It has been suggested that this is the basis for the human preference for a *flow* condition in which the challenge of a task is managed to remain between the extremes of boredom and anxiety.

In conclusion, it is suggested that the approach taken in this paper provides the foundations for a bridge between cybernetic, phenomenological models of adaptive interactions in societies of agents and knowledge-level modeling of the same phenomena. The simple model of knowledge as a state variable imputed by one agent in modeling another is sufficient to account for the performance-based training phenomena described. However, such training is powerful precisely because it is independent of exact knowledge of that state variable. If some knowledge is available then tasks may be selected that are specific to the state of knowledge of the trainee. Other knowledge transfer phenomena such as mimicry and language may also be modeled and controlled in similar terms. The higher-level knowledge processes in societies of adaptive agent may be modeled as natural extensions of the basic phenomena described.

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada.

References

Bond, A.H. and Gasser, L., Ed. (1988). **Distributed Artificial Intelligence**. San Mateo, Morgan Kaufmann.

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

Clancey, W.J. (1989). Viewing knowledge bases as qualitative models. **IEEE Expert 4**(2) 9-23.

Csikszentmihalyi, M. (1990). **Flow: The Psychology of Optimal Experience**. New York, Harper and Row.

Gaines, B.R. (1968). Training the human adaptive controller. **Proceedings of the Institution of Electrical Engineers 115**(8) 1183-1189.

Gaines, B.R. (1969). Adaptively controlled instruction for a tracking skill. **Programmed Learning Research. Sciences de Comportement**. pp.321-336. Paris, Dunod.

Gaines, B.R. (1972a). The learning of perceptual-motor skills by men and machines and its relationship to training. **Instructional Science 1**(3) 263-312.

Gaines, B.R. (1972b). Training, stability and control. **Instructional Science 3**(2) 151-176.

Gaines, B.R. (1987). Positive feedback processes underlying functional differentiation. Caudhill, M. and Butler, C., Ed. **Proceedings of IEEE First International Conference on Neural Networks. Vol.2**. pp.387-394.

Gaines, B.R. (1988). Positive feedback processes underlying the formation of expertise. **IEEE Transactions on Systems, Man & Cybernetics SMC-18**(6) 1016-1020.

Gaines, B.R. (1989). Social and cognitive processes in knowledge acquisition. **Knowledge Acquisition 1**(1) 251-280.

Gaines, B.R. (1994). The collective stance in modeling expertise in individuals and organizations. **International Journal of Expert Systems 7**(1) 21-51.

Gaines, B.R. and Andreae, J.H. (1966). A learning machine in the context of the general control problem. **Proceedings of the 3rd Congress of the International Federation for Automatic Control**. London, Butterworths.

Hoffman, D.L. and Novak, T.P. (1995). Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations. Owen Graduate School of Management, Vanderbilt University, Nashville, TN. http://www2000.ogsm.vanderbilt.edu/cmepaper.revision.july11.1995/cmepaper.html.

Langton, C.E., Ed. (1995). **Artificial Life: An Overview**. Cambridge, Massachusetts, MIT Press.

Mamdani, E.H. and Assilian, S. (1975). An experiment in linguistic synthesis with a fuzzy logic controller. **International Journal Man-Machine Studies 7**(1) 1-13.

Sugeno, M. (1985). **Industrial Applications of Fuzzy Control**. Amsterdam, North-Holland.

Zadeh, L.A. (1973). Outline of a new approach to the analysis of complex systems and decision processes. **IEEE Transactions on Systems, Man and Cybernetics SMC-3** 28-44.