POSITIVE FEEDBACK PROCESSES UNDERLYING FUNCTIONAL DIFFERENTIATION

Brian R Gaines
Department of Computer Science, University of Calgary
Calgary, Alberta, Canada T2N 1N4

Knowledge processes in neural networks are analyzed in terms of the development of functional differentiation in an initially uniform system. Socio-cultural processes are modeled by treating the biomass of the planet as a distributed network acting as an anticipatory system to maximally project its own survival. A range of positive feedback processes are sufficient to account for the development of a high degree of functional differentiation even if the overall system is initially completely undifferentiated. This model is attractive in enabling cultural and linguistic aspects of knowledge acquisition to be treated within a single parallel distributed processing framework.

1 Introduction

An attractive feature of neural networks as computational tools is the potential to take a large-scale parallel processing system of undifferentiated units and form it into a highly targeted computing engine, not by detailed interior programming, but by external environmental control. This *ecological* approach to programming is attractive not only because it bypasses the technical/conceptual barriers to programming large-scale, parallel systems, but also because the resultant engine has intrinsic robustness and adaptability that it seems impossible to achieve with conventional programmed systems. There is a *natural* relationship between the derived functionality of the network, its tasks, and the mode of derivation.

This paper analyzes the systemic principles underlying the formation of functionality in undifferentiated networks, showing that positive feedback in various forms is sufficient to account for the phenomena of differentiation. The wide variety of environmental processes that produce such feedback are then analyzed to show how the ecology of programming sources arises. Simulation results are given showing these processes in action.

It is suggested that the processes of formation of differentiated functionality in neural networks need to be considered in the wider framework of socio-cultural processes of knowledge transfer. The complete biomass of the planet may be viewed as a single neural network which is physically distributed in the brains of individuals. This allows socio-cultural processes to be viewed as intra-net phenomena that are only secondarily inter-brain phenomena because of the physical distribution of the network. In particular, linguistic phenomena may be incorporated in the model as part of the coordination processes of the distributed network.

2 The Biomass as a Single Distributed Neural Network Acting as an Anticipatory System

Many theorists have pointed to the strong analogies between different levels of biological functioning. Miller gives a comprehensive account of a wide range of living systems based on a common detailed systemic model instantiated at different levels of the bio-system [12]. He distinguishes 7 levels:

Supranational System
Society
Organization
Group
Organism
Organ
Cell

Figure 1 Miller's seven levels of biological systems

It is interesting to consider the top level system, effectively the biomass of the planet, as the basic system that acquires knowledge from experience, and regard the lower levels as the functional differentiation of the biomass. A neural network in Miller's framework is a particular *organ* made up of specialist *cells* (neurons) at the lower level, and serving the command and control functions of the *organism* at the level above. In so doing it enables the organism to play specialist roles with *groups* that themselves have specialist roles within *organizations* forming *society* and being part of the *supranational system*, or biomass.

The command and control functions of the network may be given a general, but operational, definition by considering the overall function of the biomass to maximally project its own survival by anticipating the future. That is we consider the overall system to be *anticipatory* in nature [14] and develop the logic of its subsystems in terms of this overall anticipatory goal. The primary drive of a living system is its survival. To survive, the system must have access to the necessary resources and be capable of coping with threats to its survival. Systems evolve to maximize their access to resources and minimize their vulnerability to threats. In a static universe a simple model of resource availability and prey-predator relations determines the dynamics of living systems.

The beginnings of the human race were set in a static universe but its activities soon began to change that universe so that uncertainty about the future became the predominant factor in human survival processes. Changes in the earth that would have taken millions of years began to occur over millennia. Now multiple waves of change occur within human lifetimes. Humanity developed resources far beyond their natural availability and, in so doing, changed the ecology of the earth. It extinguished all predators other than itself and again changed the face of the earth. Figure 2 shows the system dynamics underlying humanity [3]. The shaded areas show the response to uncertainty. The keys to survival in an uncertain world are three-fold: first to maximize territorial dispersion so that some part of the system is outside the range of a threat; second to maximize cultural diversity so that some part of the system has the capability to survive a threat; third to improve anticipation of the future, passively to predict threats in advance, and actively to rebuild the universe so that they do not occur—those are the roles of science and technology.

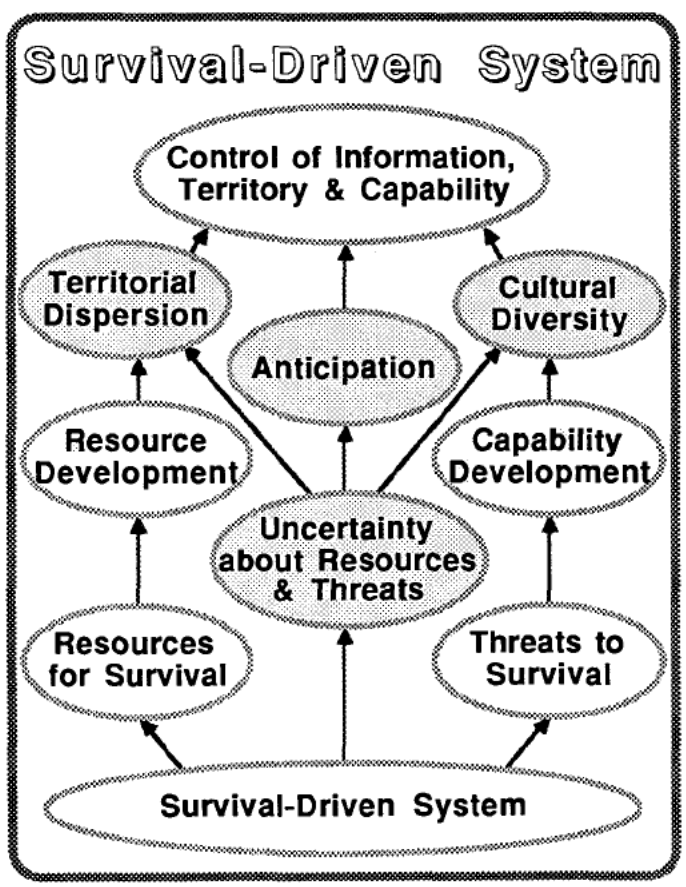

Figure 2 Humanity as a survival-driven anticipatory system

It is conventional to treat an anticipatory system as a model-forming system since that is a natural correlate of anticipation as an effective predictor or creator of the future. It is then conventional to treat the human brain as essentially a modeling system and develop cognitive psychologies upon this foundation. However, this is misleading if one then looks for the 'modeling algorithms' or the 'model.' A neural net perspective is more appropriate because it accounts for the anticipation without specific modeling algorithms or models. As Smolensky has noted, the apparently bottom-up approach commencing with the net is actually conceptually top-down in accounting for cognition [16]. Figures 1 and 2 give some intuition as to why this is so—the actual net to be considered is a global entity and its survival is its primary drive. What we see as nets at the brain, schema and neuronal levels are entities generated through the anticipatory processes of survival.

As an analogy to underpin the intuition, consider the process whereby sand takes up the shape of an irregular container. We can look at the end result and say that the sand has 'modeled' the container. However, we would find no modeling algorithm, and the existence of the model itself is a construct of our observation that the sand and the container now match one another globally—it is only in a very remote sense a property of the sand. If we wished to expedite and improve the process—to *educate* the sand—we would be foolish to attempt to aid the local *modeling* process by moving individual grains of sand. Instead we would shake the container to provide an environment of maximum opportunity for the sand to settle into the shape of the container. We would shake roughly at first and then increasingly gently.

This is exactly the *modeling* process of the Boltzmann machine and its management through simulated annealing [15]. It also corresponds to Gibson's *ecological* model of perception in which we model the world by the brain coming into equilibrium with the stream of incoming sensation [4]. The theoretical foundations for the relation between an anticipatory system and its apparent modeling capability can also be given a simple interpretation. If we consider a perfect anticipatory system that is able to match the output of the system it is anticipating symbol by symbol, then an observer modeling each of the two systems through a behavior-structure adjunction [5] will ascribe the same minimal structure to both up to an isomorphism. Because the observer obtains the same model of both systems, he may suppose that the anticipatory system contains such a model.

These concepts bridge between the amorphous uniformity of nets and the highly structured concepts we have of modeling processes. It is convenient to consider modeling structures as a means of classifying the behaviors and interactions of anticipatory systems at a phenomenonological level. However, it is a category error to assume that these structures need necessarily exist at a neurological level.

3 Modeling and the Nature of Expertise

The anticipatory system model of the distributed processes of the biomass developed above may now be used to give an account of the formation and role of expertise in society, a particularly interesting example of functional differentiation significant for human knowledge processes and their support through information technology. Hawkins has abstracted from experience at Schlumberger in developing mineral exploration expert systems and proposed a model of human expertise relevant to expert systems [8]. The expert develops a minimal model that accounts for the data provided by the user, feeding back to the user this model, some aspects of its derivation from the information given, and the practical consequences of the model in terms of advice. The user plays an active role in further developing the model by providing more information until he or she is satisfied with the model and consequent advice. Thus, advice giving and taking is a matter of negotiation around a process of model formation.

This model of the expert-user interaction as the negotiation of a mutually acceptable model may be given formal foundations in mathematical modeling theory. Figure 3 shows a general hierarchical modeling process [2] based on Klir's analysis of model formation [10]:

- *Constructs* are those distinctions that the particular modeling system makes, a language for describing of the world;
- *Data* are descriptions of actual case histories in terms of the constructs, an account of experiences of the world;
- *Hypotheses* are the means of regenerating particular case histories from generalized accounts, rationalizations of the world (often called *models*);
- *Analogies* are similarities between differently generated generalized accounts, correspondences between models;
- *Abstractions* are accounts of a wide range of models, underlying foundations of analogies;
- *Principles* are systemic foundations for abstraction, accounts of abstractions.

Modeling may be evaluated by the degree to which a level accounts for the information flowing through it. This may be measured as the *surprise* at an actual event in relation to one predicted, and the expected surprise is a form of entropy [2]. The formal theory of modeling is one in which selections are made at each level down the hierarchy to minimize the rate at which surprise is passing up the hierarchy. The criteria for selection independent

of the data are generally thought of as being ones of simplicity/complexity: of two models which fit the data equally well choose the simplest. However, notions of simplicity/complexity are not well-defined nor intrinsic to the class from which selections are made. The simplicity/complexity ordering is arbitrary and in its most general form is just one of *preference*. Hence the general modeling schema is one in which surprise flows up the hierarchy and preference flows down, signals with definite neurological counterparts [18]. In situations that are mathematically well-defined, such as determining the structure of a stochastic automaton from its behavior, such a model schema gives the correct results. Conversely, the success of the schema in stabilizing with regard to a given world defines the characteristics of that world.

It should be noted that the modeling is an active process also—constructs are not just static partitions of experience. They may be operations: actions in psychological terms; processes in computational terms. Whether a system finds a distinction in the world, imposes it passively as a view of the world, or imposes it actively as a change in the world, is irrelevant to the basic modeling theory. Figure 3 encompasses both the expert's model-formation and model-application processes.

Hawkins' studies exemplify the roles of each of these levels in expert behavior and in the negotiation between expert and seeker of advice. In particular, he illustrates the significance of shared constructs between expert and user, and the potential for 'data' from falsely constructed events to invalidate expertise. What the basic modeling theory does not account for, however, is these types of social interaction, the relation between experts and users, the assumptions that an expert makes about the problem situation, the relation between experts and society. Indeed, it does not account for the concept of an 'expert' except as a particularly well-validated modeling system, leaving open the question as to how one system becomes better validated than another.

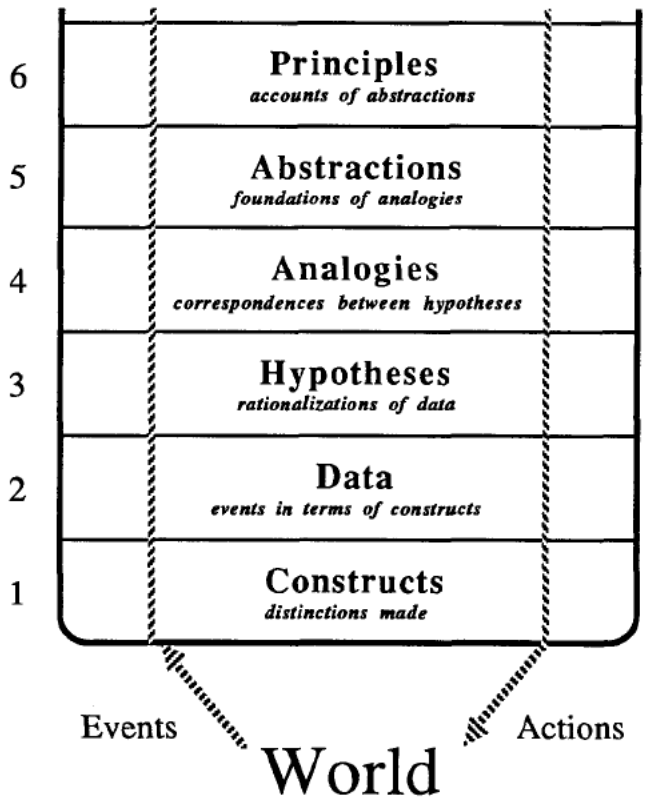

Figure 3 Levels in a general modeling system

4 Reflective Equilibria Underlying Induction and Expertise

The weakness of the account of modeling given in Figure 3 is that it accounts for only a minor part of the knowledge acquisition processes of an individual. Most of our knowledge does not come from forming models of events in the world, but rather from the social transfer of others' knowledge by a variety of cultural processes, some formalized in education but many very informal. However, the account can be extended to encompass such cultural knowledge acquisition if one notes that human civilization as a whole is representable as a general modeling system having no sources of knowledge other than its combined, cumulative experience of the world. The modeling schema of Figure 3 applies as much to groups of people, companies and societies as it does to the roles of a person. The modeling processes of an individual may be viewed as a cross-section of those of the civilization of which they are part. The modeling processes of society may be viewed as carried out by a distributed system of sub-processes in individuals. Cultural processes of knowledge transfer are then just various modes of communication between such sub-processes. The distributed system analogy nicely accounts for a variety of psychological and social phenomena: one person may assume many psychological roles (process switching), whereas a group of people working together may act as a single modeling entity and hence behave as one process (distributed processing).

These concepts may be formalized by considering the inductive inference process underlying knowledge acquisition. Deduction guarantees to take us from valid data to valid inferences, but the inferences are thereby part of the data—no new knowledge is generated. Induction takes us from valid data to models of that data that go beyond it—by predicting data we have not yet observed, and by giving explanations of the data in terms of concepts that are unobservable. Induction generates new knowledge but, as Hume pointed out over 200 years ago, the process is not deductively valid and it is a circular argument to claim that it is inductively valid [9]. Goodman proposed that we accept the circularity but note that it involves a dynamic equilibrium between data and inference rules as shown in Figure 4: "A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend" [6]. Rawls in his theory of formal justice as induction from informal moral behavior terms this a *reflective equilibrium* [13].

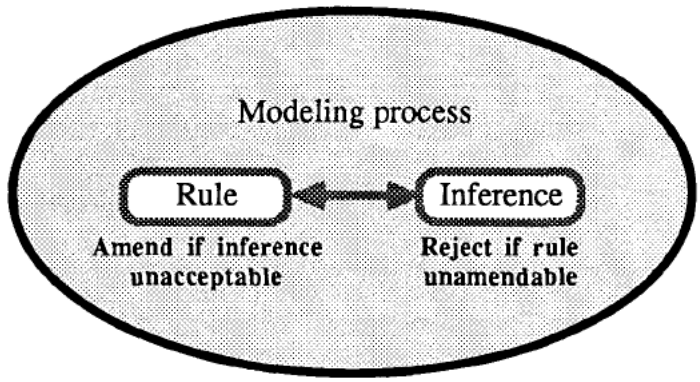

Figure 4 Reflective equilibrium in inductive inference

Stich and Nisbett noted flaws in Goodman's argument in that individuals do not necessarily behave rationally in their own behavior relative to known good practice [17]. They repaired them by proposing that the equilibrium is social not individual: "a rule of inference is justified if it captures the reflective practice not of the person using it but of the appropriate experts in our society." This gives an operational definition of the role of experts within an overall knowledge acquisition system as referential sub-systems managing a domain of the knowledge acquisition process. It also leads to an concept of the role of expert systems in society, as making the referential process more overt and widely available.

Thus the pragmatic model of expertise developed by Hawkins can be brought into the formal framework of modeling theory by taking society to be a distributed modeling system in which experts are those with responsibility for managing the inductive process in particular domains. However, even this extension does not specifically account for the formation of expertise—how does a particular individual become an expert. The following section shows that random processes in society are adequate to account for expertise formation.

5 Functional Differentiation Through Positive Feedback

An expert in Stich and Nisbett's account is a modeling system that has particular authority to decide on changes when events violate rules—in terms of Figure 3, are new constructs needed, should the data be rejected, should the hypothesis be changed, and so on. However, this authority is a social construct, not usually one of right in present society, and the ascription of such authority to an individual is part of the larger modeling system's processes in determining which of its sub-processes is effective. Now, to become effective in modeling is, from Figure 3, dependent on many factors: having access to knowledge at the upper levels such as principles, abstractions, analogies and classes of hypothesis; having access to knowledge at the lower level such as adequate construct systems and relevant data from case histories. Such access is itself a social resource, not generally available to all, and there are social mechanisms for allocating such access to individuals. It is the interaction between the criteria for access and the acquisition of knowledge that leads to the formation of expertise.

Sociologists have noted the very strong positive feedback processes in the dynamics of the scientific community [7]. Merton 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 [11]. The qualitative effect of such positive feedback processes is to amplify random differences within a population, creating strong distinctions which have no other basis than the feedback process itself. That is, even if all scientists were created equal in their capabilities, some would become experts relative to others because random

differences in early performance generated credit that affected resource allocation and gave them greater opportunities to acquire knowledge.

In many areas of expertise where expert systems are expected to play a major role these positive feedback processes are exceptionally strong. In medicine, for example, the 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, strongly funnels problems to those who are regarded as experts. It is, however, precisely these problems from which new knowledge is generated. Hence the data at level 2 in Figure 3 is made available primarily to those with already established reputations for having capabilities at level 3 and above to deal with it. Similar considerations apply to the award of scholarships, invitations to scientific congresses, and so on. There are many positive feedback loops that seek out expertise and seek to develop it. However, they are also adequate in their own right to create it.

Simple simulation experiments are adequate to show the strength of these phenomena. Figure 5 shows the results of a simulation in which a user community determines which of two individuals it will use in the solution of a succession of problems. The rules of the simulation are: equal learning capability for each individual; a finite number of items of knowledge; each problem requires certain knowledge; if the individual does not have the knowledge necessary it guesses with a probability of success, learning if it succeeds; the community choses the individual for a problem with equal probability initially, gradually biasing the choice according to success or failure; there is no communication of knowledge between individuals. Figure 5 plots: the probability that one individual will be always preferred and shows that this rapidly approaches 1.0—an 'expert' is determined; the expected knowledge of the better individual and 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 individuals becomes the 'expert' is, of course, completely chance.

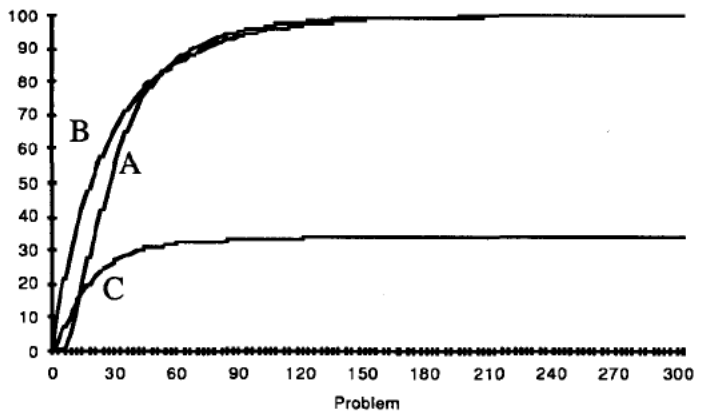

A: Percentage probability that problem will be given to only one of the two 'experts'
B: Expected knowledge of superior 'expert'
C: Expected knowledge of inferior 'expert'

Figure 5 Simulation of the formation of expertise through positive feedback

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 cultural knowledge transfer. All the simulations bear out the expected qualitative result, that a variety of different positive feedback mechanisms in society are adequate to account for differential expertise in a sample with initially equal knowledge and capabilities.

These results are not intended to suggest that individual differences are insignificant, or be a basis for egalitarian ethics. They do demonstrate, however, that the process of expertise formation is a natural systemic phenomenon in a learning society. It is functional because it distributes the processors available to different tasks through

global mechanisms with no localized control. The random mechanism is efficient because it does not matter to the overall system which processors are allocated to particular tasks, and it will tend to pick out individual differences in capabilities if they exist. It puts experts into the role of traders who must offer their services to solve problems in order to gain access to those problems in order to increase their knowledge.

In practice, the professional and scientific communities also have strong mechanisms set up to reduce the effects of positive feedback in exaggerating differences, such as the ethical pressure to disseminate knowledge as widely as possible, and mechanisms for doing this through publications and seminars. However, much knowledge underlying skilled performance cannot be shared in this way and the positive feedback mechanisms will always exert a strong effect.

6 Functional Differentiation at Lower and Higher Levels

The model of the formation of expertise through positive feedback processes leading to functional differentiation has been developed for the middle levels of the network which is the biomass. It is interesting to speculate that exactly analogous process take place at both lower and higher levels. Phylogenetically, Dawkins has given many examples of the significance of positive feedback processes in directing evolution [1].

At the lower levels of neural nets functioning within a brain the formation of schema may be seen as evidence of functional differentiation. Even if these are completely distributed across the network it is still possible to consider them as both psychological and computational differentiations. For example, one may separate part of the input to the net as a 'context' and consider the behavior of the net in response to the remaining input occurring within that context as generating a schema. There are probably also neurological structures corresponding to such contexts and schema. The attention directing processes mediated through the limbic region are a neurological mechanism for establishing such contexts.

At the higher levels of groups, societies and so on, there are socio-economic pressures leading to functional differentiation as described. It is interesting to speculate that the hierarchy goes on beyond that analyzed by Miller in Figure 1, and that the earth with its biomass may be itself a functionally differentiated unit in some arena of living systems of which we are not yet aware.

7 Conclusions

The adaption of the species to the total ecology is a holistic process best understood in terms of the feedback processes involved and the natural relations induced. A number of perspectives on the derived modeling processes may be given: that of anticipation of the physical environment through its modeling and manipulation to increase the probability of survival; that of formation of a socio-cultural infrastructure through the coordination processes of the distributed brain in this anticipatory activity; that of formation of a knowledge, or memetic, infrastructure through the by-products of time-binding in this anticipatory activity.

The theory developed may be extended to distributed nets where inter-net information flow is very much weaker than intra-net flow, and this can be treated as a two-layer net with the same systemic principles applied to the overall upper layer as applied within each net in the lower layer. That is, ecological positive feedback loops generate functional differentiation between nets now rather than within nets, and higher-level 'models' are distributed over sub-systems of nets. This extended theory accounts for the processes of knowledge acquisition within the human species, and the variety of environmental processes producing positive feedback are cultural mechanisms underlying the creation of knowledge. The extended theory accounts naturally for the *memetic* processes whereby knowledge itself has its own dynamics and ecology. The human species may be regarded as a unitary system producing knowledge as a by-product of its survival processes. It has a distributed brain composed of physically distributed neural networks having a degree of autonomy within the coupling of a cultural matrix.

This model has been used to model the role of the expert in the knowledge processes of society—that of managing the inductive dynamics of knowledge acquisition. Experts trade the skills based on their knowledge for access to problems providing them with new material for the inductive process. The formation of expertise is functional in general because it leads to division of labor in the management of knowledge acquisition. The development of an individual expert is a random process brought about by strong positive feedback loops in the social process, for example, that a proto-expert with superior performance is brought more problems and hence has a greater opportunity to learn and improve that performance. A diversity of such positive feedback processes operate in the professions and sciences with little relation between them except their overall effect in promoting the formation of expertise. Sociological evidence for the existence of these effects has discussed, and simulation studies showing the quantitative phenomena described.

8 Acknowledgements

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

9 References

[1] Dawkins, R. (1986). **The Blind Watchmaker.** Harlow, Essex, UK: Longman.

[2] Gaines, B.R. (1977). System identification, approximation and complexity. **International Journal of General Systems, 3,** 145-174.

[3] Gaines, B.R. (1986). Foundations of knowledge engineering. Bramer, M.A., Ed. **Research and Development in Expert Systems III.** pp.13-24. Cambridge: University Press.

[4] Gibson, J.J. (1979). **The Ecological Approach to Perception.** Boston: Houghton Mifflin.

[5] Goguen, J.A. (1973). Realization is universal. **Mathematical Systems Theory 6.** 359-374.

[6] Goodman, N. (1973). **Fact, Fiction and Forecast.** Indianapolis: Bobbs-Merrill.

[7] Hagstrom, W.O. (1965). **The Scientific Community.** New York: Basic Books.

[8] Hawkins, D. (1983). An analysis of expert thinking. **International Journal of Man-Machine Studies, 18**(1), 1-47 (January).

[9] Hume, D. (1739). **A Treatise of Human Nature.** London: John Noon.

[10] Klir, G.J. (1985). **Architecture of Systems Problem Solving,** New York: Plenum Press.

[11] Merton, R.K. (1973). **The Sociology of Science: Theoretical and Empirical Investigations.** Chicago: University Press.

[12] Miller, J.G. (1978). **Living Systems.** New York: McGraw Hill.

[13] Rawls, J. (1971). **A Theory of Justice.** Cambridge, Massachusetts: Harvard University Press.

[14] Rosen, R. (1985). **Anticipatory Systems.** Oxford: Pergamon.

[15] Rumelhart, D.E., McClelland, J.L. & PDP Research Group. (1986). **Parallel Distributed Processing.** Cambridge, Massachusetts: MIT Press.

[16] Smolensky, P. (1986). Formal models of subsymbolic processes: an introduction to harmony theory. Sharkey, N.E., Ed. **Advances in Cognitive Science 1.** pp. 204-235. Chichester, UK: Ellis Horwood.

[17] Stich, S.P. & Nisbett, R.E. (1984). Expertise, justification and the psychology of inductive reasoning. Haskell, T.L., Ed. **The Authority of Experts.** pp. 226-241. Bloomington, Indiana: Indiana University Press.

[18] Tucker, D.M. & Williamson, P.A. (1984). Asymmetric neural control systems in human self-regulation. **Psychological Review, 91**(2), 185-215 (April).