Journal of Intelligent Information Systems 9, 277–298 (1997)
© 1997 Kluwer Academic Publishers. Manufactured in The Netherlands.

Knowledge Management in Societies of Intelligent Adaptive Agents

BRIAN R. GAINES gaines@cpsc.ucalgary.ca
Knowledge Science Institute, University of Calgary, Alberta, Canada T2N 1N4; http://ksi.cpsc.ucalgary.ca/KSI

Abstract. A model is developed of the emergence of the knowledge level in a society of agents where agents model and manage other agents as resources, and manage the learning of other agents to develop such resources. It is argued that any persistent system that actively creates the conditions for its persistence is appropriately modeled in terms of the rational teleological models that Newell defines as characterizing the knowledge level. The need to distribute tasks in agent societies motivates such modeling, and it is shown that if there is a rich order relationship of difficulty on tasks that is reasonably independent of agents then it is efficient to model agents competencies in terms of their possessing knowledge. 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 provides a basis for assigning a granularity to knowledge that enables learning processes to be managed simply and efficiently.

Keywords: knowledge management, knowledge level, intelligent agents, adaptive agents, management of learning, learning curves

1. Introduction

The past decade has seen a major growth in research on the structure, behavior and applications of systems based on communities of intelligent adaptive agent. This research has been undertaken under many different auspices, as research in distributed artificial intelligence (Gasser and Avouris, 1992), artificial life (Langton, 1995), intelligent agents (Wooldridge et al., 1996), reactive agents (Castelfranchi and Müller, 1995), multi-agent systems (Weiß and Sen, 1996), reasoning about knowledge (Fagin et al., 1995), and so on. The artificial intelligence literature has focused on issues relating to the design of computational artifacts, but it has been paralleled by a closely related literature that examines the phenomena of natural communities of intelligent adaptive agents from perspectives such as cultural biology (Boyd and Richerson, 1985), and the active, dynamic or network society (Etzioni, 1968; Castells, 1996; Snooks, 1996). There has also been a growth in recent years of organizational literature on the management of knowledge in companies modeled as communities of intelligent adaptive agents (Barabba, 1995; Nonaka and Takeuchi, 1995; Johansson and Nonaka, 1996).

What these literatures have in common is a focus on the emergence of knowledge processes through interactions in communities of intelligent adaptive agents. They situate knowledge not in the operation of a particular agent but rather in the social processes whereby agents coordinate their behaviors. However, much of the literature takes knowledge itself as a starting point, a given entity that is part of the notion of an intelligent agent, and focuses on

knowledge acquisition, inference and communication. The question of how the behaviors of agents in communities emerge from the situational dynamics in such a way that they can be imputed goals and knowledge is of increasing interest in the artificial intelligence literature (Castelfranchi and Müller, 1995; Maes, 1995), as is the management of the processes of making overt "tacit knowledge" in the organizational literature (Nonaka and Takeuchi, 1995). However, there is as yet no overall theory of the nature of the knowledge construct in terms of its emergence through the social processes of communities of intelligent adaptive agents.

The objective of the research described in this paper is to derive fundamental principles for knowledge management in societies of intelligent adaptive agents by defining knowledge in operational terms and using this definition to analyze the knowledge dynamics of organizations composed of agents and technologies.

First, Newell's (1982) *knowledge level* studies are recapitulated to show that knowledge can be treated as a state variable imputed to an agent by a modeler to account for its behavior.

Second, his rational teleological model is shown to involve few fundamental presuppositions about the nature of the systems involved, except that they persist in time and the observer hypothesizes that they actively bring about this state of affairs.

Third, it is shown that the colloquial interpretation of knowledge as something material possessed by an agent arises naturally in accounting for the capabilities of agents to perform tasks.

Fourth, it is shown that further constraints on knowledge level modeling arise from the hypothesis of compositionality in the derivation of the capabilities of a team from its component agents. Furthermore, this give rise to knowledge modeling of organizational knowledge, including that of supporting technologies.

Fifth, it is shown that further constraints on knowledge level modeling arise from the hypothesis that an agent's learning can be managed through the regulation of a graded sequence of tasks that it is given to perform.

In conclusion, it is suggested that a knowledge level analysis of agents, organizations and technologies provides appropriate formal foundations for knowledge management in societies of intelligent adaptive agents.

2. The knowledge level

In his seminal paper on the *knowledge level* Newell (1982) situates knowledge in the epistemological processes of an observer attempting to model the behavior of another agent:

> "*The observer treats the agent as a system at the knowledge level, i.e., ascribes knowledge and goals to it.*" (p. 106)

emphasizing that:

> "*The knowledge level permits predicting and understanding behavior without having an operational model of the processing that is actually being done by the agent.*" (p. 108)

He defines knowledge as:

> "*Whatever can be ascribed to an agent such that its behavior can be computed according to the principle of rationality.*" (p. 105)

noting that:

> "*Knowledge is that which makes the principle of rationality work as a law of behavior.*" (p. 125)

and defining rationality in terms of the principle that:

> "*If an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action.*" (p. 102)

Newell's argument form is a *cybernetic* one of the type originated by Wiener (1948) and refined by Ashby (1956) whereby an arbitrary system is treated as a *black box* to be modeled on the basis of its input/output behavior with no presuppositions about its internal structure. Ashby (1952) used this argument form to derive many phenomena of living systems, such as habituation, from general properties, such as the existence of many alternative attractors in the state system. Zadeh (1964) developed the abstract formulation of system identification from a cybernetic stance, showing how the notion of *state* is an abstraction introduced in modeling formalisms to account for the influence of past experience on future behavior. Gaines (1977) developed general algorithms for such identification in terms of arbitrary measures of model *complexity* and of the *approximation* of a model to observed behavior, and showed that appropriate measures led to optimal identification of deterministic and stochastic automata from their behavior. He emphasizes the formal arbitrariness of the presuppositions underlying a modeling schema, and shows that inappropriate presuppositions lead to indefinitely complex models (Gaines, 1976).

In the light of these analyses, Newell's arguments may be seen as stating that knowledge is a state variable imputed by a modeler in order to account for its behavior, and that the appropriate presuppositions for modeling an agent are those of rational teleology, that it has goals and acts to achieve them. Two fundamental questions arise about Newell's framework for knowledge, one reaching backwards to the justification of modeling behavior teleologically in terms of goals and their rational achievement, and the other reaching forwards to the nature of the knowledge state space that an observer will generate, its detailed qualitative and quantitative characteristics.

The next section briefly examines the preconditions for rational teleological models to be effective, and the remainder of the paper develops in depth the structure of knowledge models that will arise in a society of agents.

3. Emergence of rational teleological models

One way of analyzing the foundations of rational teleological models is to assume that they have none—that the modeling of other agents in terms of goals and knowledge is

justified to the extent that it works—a pragmatic argument of the form developed by Peirce and James (Ayer, 1968). This assumption is that of Dennett's (1987) *intentional stance*, and it is in accordance with the basic theory of modeling, the Popperian position that our presuppositions in modeling are but *conjectures* subject to *refutation* if we are not satisfied with the results of using them (Popper, 1963). Modeling theory tells us that if the intentional stance was not appropriate to modeling human agents then it would lead to complex models with poor predictive power and we would find it more useful to adopt some other stance.

However, it is useful to examine some simple systemic characteristics of agents that would justify the use of rational teleological models if only to illustrate how few presuppositions are necessary for the model to be useful (Gaines, 1994).

3.1. Rational teleological models of persistent systems

The most fundamental properties which we impute to any system are its *existence* and *persistence* over time. A system is identifiable as not having existed before some time, of definitely existing after some later time, of persisting in existence until some later time, and of not existing again after some later time. This coming into existence, persisting for while, and going out of existence again is a common property of all systems. It applies to both living and non-living systems, and in living systems it applies at all levels from cell to species.

What characterizes living systems are the recursive activities of self-replication underlying their persistence, that they actively and continually *create* the conditions for their persistence. Maturana (1975) has proposed that this is the fundamental distinction between living and non-living systems. Autopoietic systems:

> "*are systems that are defined as unities as networks of production of components that* (1) *recursively*, *through their interactions*, *generate and realize the network that produces them*; *and* (2) *constitute in the space in which they exist*, *the boundaries of this network as components that participate in the realization of the network. . . a living system is an autopoietic system in physical space.*" (Maturana, 1981)

However, there is no notion of goals or knowledge in Maturana's definition, and no ascription of intentions to living systems. A reactive persistent system in itself has no goals or intentions. It reacts to its environment through mechanisms that tend to maintain its persistence despite changes in its environment. An external observer may model this behavior as goal-directed because that provides a simple predictive explanation. That is, if an autopoietic system when disturbed, regardless of what state it is triggered into, seems to return to its original state, it is naturally modeled as goal-seeking. If the system's environment happens to contain other systems like itself and the system's activities include observation and modeling, it may model the other systems as goal-directed, and then by analogy come to model itself as goal-directed. This is a natural outcome of autopoiesis in a social environment.

As well as not reading too much into models of autopoietic systems, it is important to note that we can ascribe very little to their *existence*. A chaotic universe has a probability

of producing any system including autopoietic systems. Once such systems exist and are modeled properties emerge (Sharif, 1978). As Peirce remarks:

> "*Law begets law*; *and chance begets chance*. . . *the first germ of law was an entity which itself arose by chance*, *that is as a First.*" (Peirce, 1898)

Jantsch (1980) and Prigogine (1984) have developed detailed models of how organization emerges from chaos. Gould (1989) has analyzed the fossil record and modeled the genesis and extinction of a wide variety of species as low probability random events. Monod (1972) has given a biochemical model of life as an improbable phenomena that, once it exists, follows deterministic laws. When a living system comes into existence it acts to persist, but, from the systemic perspective advanced by Maturana, this is the definitional property by which we recognize its existence as a living system, not an additional property going beyond active persistence.

Barrow and Tipler (1986) have analyzed the remarkably narrow physical conditions under which life as we know it can exist, and when one examines the mechanisms by which a living organism narrows these conditions even further in order to persist it is natural to ascribe purpose to its activity. For example, Cannon (1932) terms such activity homeostasis and part of *The Wisdom of the Body* and Ashby (1952) in analyzing homeostasis as part of his *Design for a Brain* models it as a goal-directed process. However, he also shows how such apparently goal-directed behavior arises in any system with many states of equilibrium. The utility of an intentional stance stems from simple systemic considerations, and one has to be careful in reifying the notion of agency to realize that the additional assumption of the existence of some reified 'agent' is also a matter of utility, not of existential proof or necessity.

In Ashby's day a system that reacted to its environment by acting until it arrived in a new mode of equilibrium would be seen as not only counter-acting the effects of the environment but also arriving at some state that was determined by those effects, that is, apparently targeted upon them. Nowadays, with the realization that *strange attractors* are prevalent in all forms of physical system (Ruelle, 1989), and particularly in biological processes and their higher-order manifestations such as brains (Basar, 1990), societies (Dendrinos and Sonis, 1990) and cultural phenomena (Hayles, 1991), it would be realized that the final state may be one of very many that have the equilibrating effect but is neither determined by the effects of the environment nor targeted upon them.

In particular, the definition of *fitness* of a species in evolutionary terms is merely a restatement of the species' persistence in terms of the environment in which it persists. As Ollason argues:

> "*Biologists use the concept of fitness as the explanation of the truly inexplicable. The process of evolution is exactly what the etymology of the word implies*: *it is an unfolding*, *an indeterminate*, *and in principle*, *inexplicable unfolding.*" (Ollason, 1991)

A species is fit to exist in an environment in which it happens to persist. As noted in the previous paragraph, this does not mean it was targeted on that environment or that there is a determinate relation between the nature of the environment and the species that happens

to have evolved. The environment acts as a filter of species and those that persist are fit to survive. There are no teleological implications, and this model does not give 'survival-directed' behavior any greater probability of leading to persistence than any other behavior. Gould (1989) details the random phenomena that have made particular species fit to persist for a while in the fossil record. Bickerton (1990) argues that there is no evidence for what we deem to be high-level human traits to have survival value—intelligence and language have at least as many disadvantages as advantages, and may be seen as of negative value to the survival of the human species.

3.2. Emergence in knowledge management

This conceptual framework, emphasizing opportunistic rather than goal-directed behavior, is already a major component of the knowledge management literature. One of Bridges' recommendations in *Managing Transitions* is to *Let Go of Outcomes*:

> "*we cannot ultimately control outcomes, and when we try to, we either alienate others or drive ourselves crazy.*" (Bridges, 1991)

Johansson and Nonaka use related criteria to differentiate Western and Japanese companies in their approach to marketing:

> "*Whereas strategic planning in the West typically cascades down in logical steps from broad mission statements to more specific objectives to the enumeration of tasks, the assignment of responsibilities and the fixing of a time schedule, the Japanese approach is fuzzier. The intuitive incrementalism of the Japanese means essentially experience-based learning, a natural or 'organic' process.*" (Johansson and Nonaka, 1996)

Barabba's (1995) introductory chapter in *Meeting of the Minds* is entitled *The Late Great Age of Command and Control* and critiques the normative approach to business based on predefined objectives rather than an adaptive one based on learning from the market place, the organization's natural environment.

There is an interesting parallel on this emphasis on openness to experience in Gadamer's discussion of what it is to be an expert:

> "*The nature of experience is conceived in terms of that which goes beyond it; for experience can never be science. It is in absolute antithesis to knowledge and to that kind of instruction that follows from general or theoretical knowledge. The truth of experience always contains an orientation towards new experience. That is why a person who is called 'expert' has become such not only through experiences, but is also open to new experiences. The perfection of his experience, the perfect form of what we call 'expert', does not consist in the fact that someone already knows everything and knows better than anyone else. Rather, the expert person proves to be, on the contrary, someone who is radically undogmatic; who, because of the many experiences he has had and the knowledge he draws from them is particularly equipped to have new experiences and learn from them.*" (Gadamer, 1972)

One can paraphrase the knowledge management texts cited above as stating that an 'expert organization' is one that satisfies Gadamer's notion of what it is to be an expert person. However, it is important to note that he contrasts knowledge and expertise. While a rational teleological model may naturally emerge when modeling a persistent agent as actively involved in ensuring its persistence, the knowledge imputed is a by-product of the modeling process not the cause of the persistence. Modeling the openness and adaptivity of expertise involves multiple levels of modeling, and the observer has to introduce notions of 'meta-knowledge' or 'deep knowledge' in order to account for the processes whereby the knowledge imputed to account for specific short-term behavior changes through experience.

In particular, the notion that an expert possesses 'tacit knowledge' is misleading since 'making the knowledge overt' by modeling the expert will generate a knowledge-level model of some aspects of the expert's behavior but there is no assurance that such a model captures the essence of the expertise. It is in the interplay of human resource management, the technological support of expertise, and the essential community processes of an organization that the true subtleties of real-life 'knowledge management' arise. Any realistic model of knowledge processes in communities of intelligence, adaptive agents must account for the intrinsic uncertainties of the modeling processes involved.

3.3. Summary and implications

In conclusion, in adopting an intentional stance one is selecting a modeling schema for its simplicity, convenience and utility. Newell's notions of rationality, goals and knowledge have no epistemological content and are circularly derivable from one another as definitions of what it is to adopt an intentional stance. The knowledge level can be reified only through our first being satisfied that it has predictive capabilities, and then through our further presupposing that there must be some real phenomenon out there that makes that prediction possible.

We have to be very careful in testing both of these conditions: the reflexivity of social interactions means that changes in our behavior based on assumptions about another's intentions may lead to contingent behavior on the part of the other (Levis, 1977) giving rise to apparent predictive validity; and the predictive capabilities of a cybernetic model of a black box place very few constraints on what structure actually exists within the box.

We also have to distinguish those aspects of an agent's behavior that an observer is attempting to model. For example, modeling the agent's current skills, it capability to use those skills in specific contexts such as in a team, and its capabilities to learn to improve its skills, are three different modeling requirements that place different constraints on knowledge level modeling. The following three sections investigates each of these requirements in turn.

4. Knowledge as an imputed state variable

The previous section having warned against reading too much into knowledge level models, the current one will build such a model based on a sequence of plausible assumptions largely concerned with *cognitive ergonomics*—of building models that require as little effort to develop as possible. The starting point is Newell's notion that the knowledge level originates

in one agent attempting to model another, and hence is essentially a product of a social process. One can ask the question "why should it be valuable to model another agent" and come to the conclusion that the human species is characterized by its social dependencies, the divisions of labor whereby many of the goals of one agent are satisfied through the behaviors of others. In these circumstances one agent will model another in instrumental terms, in terms of its capabilities to carry out tasks that will lead to the modeling agent's goals being satisfied—and, vice versa, the other agent will model the first in a reciprocal fashion.

4.1. Modeling agents' capabilities to perform tasks

Consider a set of agents, A, and a set of tasks, T, such that it is possible to decide for each agent, $a \in A$ whether it can carry out a task $t \in T$. Assume, without loss of generality, that this is a binary decision in that performance at different levels is assumed to define different tasks, and that we can write $a \circ t$ for the truth value that agent a can carry out task t. We can then characterize an agent's competence, $C(a)$, by the set of tasks which it can carry out:

$$C(a) \equiv \{t \in T : a \circ t\} \tag{1}$$

If one agent knows $C(a)$ for another agent, a, then it knows its competence in terms of the tasks it can carry out and can plan to manage its goals by allocating appropriate tasks to the other agent.

However, keeping track of the competencies of relevant agents in terms of extensive sets of tasks for which they are competent is inefficient both in knowledge acquisition and storage if there are many dependencies between tasks such that the capability to carry out one task is a good predictor of the capability to carry out another. A partial order of *difficulty* on tasks, $\geq$, may be defined such that the capability to carry out a task of a given difficulty indicates the capability to carry out tasks of lesser difficulty in the partial order:

$$\forall t, u \in T, t \geq u \equiv \forall a \in A, a \circ t \rightarrow a \circ u \tag{2}$$

If, there is a rich partial order on tasks independent of agents then it becomes reasonable to attempt to represent the partial order as one embedded in the free lattice generated by some set, K, which we shall term *knowledge*. Since the free lattice generated by a set of cardinality, n, has 2^n distinct members, there is potentially an exponential decrease in the amount of information to acquire and store about an agent if it is characterized in terms of its set of knowledge rather than the set of tasks it can perform. This decrease will be realized to the extent that the embedding of the task dependencies in a free lattice involves tasks corresponding to all elements of the lattice.

Thus, we posit a set of knowledge, K, such that a task, t, is characterized by the set of knowledge, $K(t)$, required to carry it out, and the order relation between tasks corresponds to subset inclusion of knowledge:

$$\forall t, u \in T, t \geq u \equiv K(t) \supset K(u) \tag{3}$$

An agent, a, is characterized by the knowledge it possesses, $K(a)$, and this determines its competence in terms of tasks:

$$C(a) \equiv \{t \in T : K(a) \supset K(t)\} \tag{4}$$

The development to this stage parallels that of knowledge spaces as defined by Falmagne et al. (1990), and applied by them to testing a student's knowledge. However, the move from an extensional specification in terms of tasks to an extensional specification in terms of knowledge is inadequate to account for situations where the capability to carry out one task may indicate the capability to carry out an infinite number of lesser tasks. Extensionally, this involves indexing the lesser tasks as involving an infinite set of knowledge, but, as Newell (1982) notes it is better represented by a schema in which knowledge is *generated* from knowledge.

If x is a subset of knowledge then $G(x)$ may be defined as the subset which can be generated from it subject to the obvious constraints that:

— the original knowledge is retained:

$$x \subset K \rightarrow x \subset G(x) \tag{5}$$

— all of the knowledge that can be generated is included:

$$x \subset K \rightarrow G(G(x)) \subset G(x) \tag{6}$$

— additional knowledge generates additional knowledge:

$$x \subset y \subset K \rightarrow G(x) \subset G(y) \tag{7}$$

Tarski (1930) noted that the consequence operator of any deductive system has these properties, and Wójcicki (1988) has used it conversely to characterize any closure operator satisfying (5) through (6) as a *logic*. As Rasiowa and Sikorski (1970) remark:

> "*the consequence operation in a formalized theory T should also be called the logic of T*"

that is, every generator defines a formal logic.

The closure operator makes it natural to stratify a set of knowledge into distinct subsets closed relative to the closure operator such that some unions of the subsets are open. That is, combining knowledge from certain pairs of the stratified subsets allows addional knowledge to be inferred. This is the basis for Klir's (1976, 1985) epistemological hierarchy of knowledge levels in modeling that can be used to model processes of knowledge transfer between agents in terms of the level of transfer (Gaines, 1989, 1994).

The development of this section has arrived at a characterization of the knowledge level that corresponds to the folk psychology notion that agents can be modeled as possessing something termed knowledge, and the cognitive science notion that the capability to generate

knowledge from knowledge corresponds to a formal deductive logic. What presuppositions have been involved in this development?

- P1: Agents have a reasonably stable competence such that it is well-defined whether an agent can carry out a task.
- P2: There is a rich order relationship of difficulty on tasks that is reasonably independent of particular agents.
- P3: Knowledge is not just possessed but can be generated in a principled fashion from other knowledge.

These are strong presuppositions but ones that seem to work reasonably well in characterizing human agents—we are acutely aware of the exceptions and treat them as anomalies.

4.2. Organizational knowledge

The previous section constrains knowledge level models to be predictive of individual agent's performance of tasks. However, agents generally work together in organizations and it is reasonable to suppose that a further constraint upon such models is that they should be predictive of the aggregate capabilities of agents operating in organizations and in conjunction with technological support.

A useful perspective from which to examine organizations is a *collective stance* (Gaines, 1994) in which humanity 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).

It is reasonable to add a further constraint to the generative function G, that:

— an agent's knowledge includes that of its components:

$$a \in b \rightarrow G(K(a)) \subset G(K(b)) \tag{8}$$

A stronger constraint may be stated as a compositional hypothesis:

- P4: The knowledge of a compound agent can be derived from the knowledge of its parts and their organization.

In practice, this may be an irrefutable hypothesis whereby we assume that if such a derivation is incorrect it is due to inadequate characterization of the agents' knowledge or of the way in which they are organized. For example, if we put together a team of people with apparently adequate knowledge between them to perform a task, and they can not do so, then we are likely to say that they lacked the skills to work together or that the situation did not allow them to. That is, we ascribe the failure of compositionality to a failure to have properly modeled the knowledge required or to an inadequate organization. We re-evaluate the knowledge model of the agents rather than ascribe the problem to a failure of

the compositionality hypothesis, thus making it an axiomatic constraint upon the notion of a complete knowledge model.

One interesting possibility is to extend the notion of knowledge to the organizational aspects of a compound agent by assessing the difference between the knowledge of an agent and that of its components, and ascribing this to its organization:

$$O(a) \equiv G(K(a)) - \bigcup_{x \in a} G(K(x)) \tag{9}$$

That is, $O(a)$ is the additional knowledge resulting from the organization of the agents into an organization.

4.3. Impact of technology on knowledge

The measurement of the impact of organizing agents in Eq. (9) may be generalized to apply to any contextual variables that impact the capabilities of an agent or agents. For example, an agent, a, together with a book, a tool, or computer support, may be regarded as an enhanced agent a', and one may measure the enhancement at the knowledge level as:

$$E(a, a') \equiv G(K(a')) - G(K(a)) \tag{10}$$

That is, $E(a, a')$ is the additional knowledge resulting from the book, tool, computer support or other contextual variables.

This analysis may be applied to give an instrumental view of the effect of one agent collaborating with another. For example, that when I help you then I am an instrument contributing to your capability. This is the form of analysis we use to explicate the notion of a coach.

One can derive a relationship between Eqs. (9) and (10):

$$O(a) = \bigcup_{x \in a} E(x, a) \tag{11}$$

That is, the organizational knowledge is the union of the enhancements that each agent contributes to the organization.

4.4. Summary and implications

This section has developed a detailed model of the knowledge level and its properties based on the notion that agents model other agents as resources capable of carrying out tasks. What the development does not do is characterize the nature of knowledge, other than as an arbitrary index set used in modeling. It would be reasonable to suppose that our actual definitions of knowledge elements would be closely related to our definitions of, and terminology for, tasks—for example, that someone capable of adding numbers might be said to have "knowledge of addition". However, too close a link to tasks would reduce the benefits of representing capabilities through subsets of knowledge rather than subsets

of tasks, and hence we would expect an attempt to characterize knowledge in a way that abstracts away from tasks and looks for more general knowledge elements that underlie a range of tasks.

However, P1 above is counter to another major reason why one agent may wish to characterize the capabilities of another. Human agents are not created with innate knowledge but instead *learn* to undertake an ever increasing set of tasks. Agents expect $C(a)$ to increase with experience and they manage the learning environments of other agents so as to maximize the rate of increase in directions appropriate to the needs of society. Thus, there is another major aspect to the characterization of an agent's knowledge—"how does it relate to the management of their learning—their training and education?"

The following sections develop a theory of the management of learning which gives further insights into the properties of a useful characterization of the knowledge level, particularly the granularity of knowledge.

5. Learning

Presupposition P1 in Section 4.1, that agents have reasonably stable competence is in contradiction to our expectations that agents will improve their competence with experience. In particular, it is antithetical to Gadamer's notion of an expert as one who learns from experience.

To take into account learning, one can weaken P1 to the presupposition that agents do not lose competence:

- P1′: Agents have a monotonically increasing competence such that if an agent can carry out a task it will always be able to carry out that task.

and treat $C(a)$ as lower bound on an agent's competence. The knowledge level analysis of Section 3 then follows but with the set of knowledge characterizing the agent's state being a lower bound on the agent's knowledge.

The analysis of Section 4 of the enhancement brought about by some supporting system may then be applied to the state of the agent: before support, with support, and after support. The after support enhancement defines the learning brought about by the experience of having the support. For example, what tasks one can perform before reading a book, while having access to it, and after having read it.

This analysis provides a basis for modeling various forms of knowledge transfer mechanisms, differentiating them from knowledge support systems (Gaines, 1990). It also making it clear that what at one time was termed 'expertise transfer' is better termed 'knowledge transfer', and that the 'transfer' is not explicit, but rather a way of describing and agent's change of state.

It is tempting to apply knowledge level analysis to the learning capabilities of an agent through the presupposition:

- P5: The capability of an agent to become competent to carry out a task can be predicted from its learning knowledge.

The notion of meta-knowledge that is predictive of an agent's capabilities to acquire knowledge is consistent with the educational literature on study skills and learning to learn (Novak and Gowin, 1984; Borenstein and Radman, 1985). However, a knowledge level analysis of learning is incomplete in that it cannot account for many of the phenomena of learning and training such as the probabilistic nature of trial and error learning and the management of training through performance feedback with no model of an agent's exact knowledge state (Gaines, 1972a, 1972b).

5.1. Uncertainty at the knowledge level

There is a fundamental uncertainty at the knowledge level in distinguishing between phenomena ascribable to the incompleteness of a model of an agent and those ascribable to the agent's learning. If the agent is in the situation of undertaking a new task and proves capable of performing it then we can ascribe this either to the agent's existing knowledge that had not been modeled or to the agent having acquired the knowledge in attempting to perform the task.

As already noted, similar considerations apply to predictions of the capabilities of a team from models of the knowledge of the agents forming the team. A full treatment of the knowledge level has to take into account that the modeling is subject to intrinsic uncertainties and that the modeled system is subject to change with experience.

This uncertainty leads to a knowledge management perspective whereby the capabilities of an agent, such as an organization, must be managed as part of the modeling process. Knowledge modeling is an active process of creating a model through action as much as it is one of fitting a model through observation.

The practical question then becomes one of how good a model needs to be for effective management. That is, we are not concerned with completeness of models but only their adequacy for particular purposes.

5.2. Coding, priming and training—The management of learning

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.

5.3. Positive feedback leads to functional differentiation

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

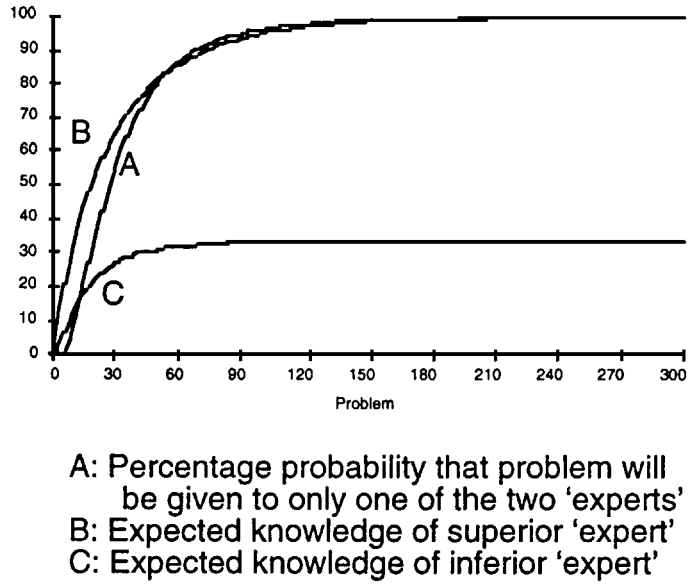

Figure 1. Simulation of the effects of positive feedback on the formation of expertise.

that one agent will be always preferred and shows that this rapidly approaches 1.0—a best 'expert' is determined. Graph B shows the expected knowledge of the better expert and graph C 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.

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'.

5.4. 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 agent's 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 again the 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 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:

$$P(n) = \sum_{0}^{d} \binom{d}{i} p^{d-i}(1-p)^{i} q^{i} = (p + (1-p)q)^{d} \tag{12}$$

and the expected rate of learning is:

$$L(n) = \sum_{0}^{d} \binom{d}{i} i p^{d-i}(1-p)^{i} q^{i} = d(1-p)\, q\, (p + (1-p)q)^{d-1} \tag{13}$$

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

$$d_{\text{opt}} = -1/\ln(p + (1 - p)q) \tag{14}$$

and the expected performance to achieve this is:

$$P_{\text{opt}} = 1/e \approx .37 \tag{15}$$

and the expected maximum rate of learning is:

$$L_{\text{opt}} = -(1 - p)q/(e(p + (1 - p)q)\ln(p + (1 - p)q)) \tag{16}$$

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

$$\underset{p\to 1}{L_{\text{opt}}} = q/e(1 - q) \tag{17}$$

The implication of Eq. (15) 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 Eq. (15) 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 with tasks of higher difficulty, and a greater chance of learning less with tasks of lower difficulty.

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 too 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 Eq. (13) for fixed n compared with the results of integrating (16), 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.

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.

Returning to the model of the knowledge level developed in Section 4, it is reasonable to hypothesize that if our knowledge level models are effective in terms of training, such

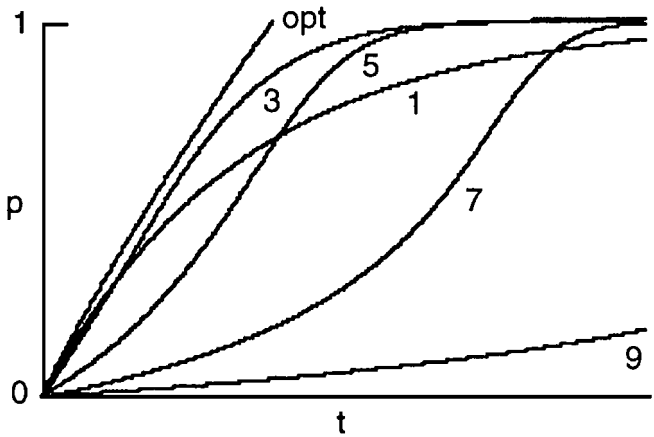

Figure 2. Learning curves for adaptive agent under various training conditions.

that optimum training leads to linear increase in knowledge with time then the natural units of knowledge are those of improvements per unit time in adaptive behavior. That is, the *granularity* of knowledge should tend to be such that the feedback training strategy is effective in linearizing the learning curve. This is the transition that society undergoes in new areas of knowledge between the era of the "inventor" when the learning curves are long and sigmoidal, to the era of the "textbook" when learning curves are comparatively short and linear. The transition is one of developing knowledge structures that make learning manageable.

5.5. A corollary: 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 5.5, 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.

6. Conclusions

In a society of agents the knowledge 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.

An operational definition of *knowledge* has been developed through a knowledge level analysis of the emergence of the knowledge construct through modeling and management processes in societies of adaptive agents. The model for the knowledge level is based upon the notion that knowledge is a state variable imputed by one agent to another in modeling its capabilities to carry out tasks. This is a cognitively efficient model if there is a rich order relationship of difficulty on tasks that is reasonably independent of agents and enables competence on one task to be predicted from that on others. The analysis shows how knowledge becomes ascribed to agents, organizations and technologies, and how formal logics of knowledge, organizations and technologies emerge naturally from reasonable presuppositions in the modeling process.

Task allocation between agents results in *functional differentiation* in an initially homogeneous society. 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.

Intrinsic uncertainties in our models and the capabilities of agents to learn imply that knowledge modeling has to be an active process of *knowledge management*. We are as much creating a model through action as fitting a model through observation. What has happened in recent years is the recognition that processes of organizational management, including personnel selection and placement, career development, team building, and so on, may all be subsumed within a single framework by knowledge level analysis of the organization as an agent with certain capabilities.

The analysis of organization as agents requires existing knowledge level theories to be extended to take into account the relations between an agent and its parts, including other agents and technologies. It also requires the theories to be extended to take into account the uncertainties in models and the learning capabilities of agents. This article provides a preliminary account of such extensions as a first step towards principled foundations for knowledge management.

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

Ashby, W.R. (1952). *Design for a Brain*, London, UK: Chapman & Hall.

Ashby, W.R. (1956). *An Introduction to Cybernetics*, London, UK: Chapman & Hall.

Ayer, A.J. (1968). *The Origins of Pragmatism*, London: MacMillan.

Barabba, V.P. (1995). *Meeting of the Minds: Creating the Market-Based Enterprise*, Boston, MA: Harvard Business School Press.

Barrow, J.D. and Tipler, F.J. (1986). *The Anthropic Cosmological Principle*, Oxford: Clarendon Press.

Basar, E. (Ed.) (1990). *Chaos in Brain Function*, Berlin: Springer.

Bickerton, D. (1990). *Language and Species*, Chicago: University of Chicago Press.

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

Borenstein, S. and Radman, Z.A. (1985). *Learning to Learn: An Approach to Study Skills*, Dubuque: Kendall/Hunt.

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

Bridges, W. (1991). *Managing Transitions: Making the Most of Change*, Reading, MA: Addison-Wesley.

Cannon, W.B. (1932). *The Wisdom of the Body*, London.

Castelfranchi, C. and Müller, J.-P. (Ed.) (1995). *From Reaction to Cognition*, Berlin: Springer.

Castells, M. (1996). *The Rise of the Network Society*, Oxford: Blackwell.

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

Dendrinos, D.S. and Sonis, M. (1990). *Chaos and Socio-Spatial Dynamics*, Berlin: Springer.

Dennett, D.C. (1987). *The Intentional Stance*, Cambridge, MA: MIT Press.

Etzioni, A. (1968). *The Active Society: A Theory of Societal and Political Processes*, New York: Free Press.

Fagin, R., Halpern, J.Y., Moses, Y., and Vardi, M.Y. (1995). *Reasoning about Knowledge*, Cambridge, MA: MIT Press.

Falmagne, J.-C., Koppen, M., Villano, M., Doignon, J.-P., and Johannesen, L. (1990). Introduction to Knowledge Spaces: How to Build, Test and Search Them. *Psychological Review*, 97(2), 201–224.

Gadamer, H.G. (1972). *Wahrheit und Methode*, Tübingen: Mohr.

Gaines, B.R. (1968). Training the human adaptive controller. *Proceedings of the Institution of Electrical Engineers* (pp. 1183–1189), Vol. 115, No. 8.

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. (1976). On the Complexity of Causal Models. *IEEE Transactions on Systems, Man and Cybernetics SMC*, 6(1), 56–59.

Gaines, B.R. (1977). System Identification, Approximation and Complexity. *International Journal of General Systems*, 2(3), 241–258.

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

Gaines, B.R. (1988). Positive Feedback Processes Underlying the Formation of Expertise. *IEEE Transactions on Systems, Man and 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. (1990). Knowledge Support Systems. *Knowledge-Based Systems*, 3(3), 192–203.

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.

Gasser, L.G. and Avouris, N.M. (Ed.) (1992). *Distributed Artificial Intelligence: Theory and Practice*, Amsterdam: Kluwer.

Gould, S.J. (1989). *Wonderful Life: The Burgess Shale and the Nature of History*, New York: Norton.

Hayles, N.K. (Ed.) (1991). *Chaos and Order: Complex Dynamics in Literature and Science*, Chicago: University of Chicago Press.

Hoffman, D.L. and Novak, T.P. (1995). Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations. *Journal of Marking*, to appear (July).

Jantsch, E. (1980). *The Self-Organizing Universe*, Oxford: Pergamon Press.

Johansson, J.K. and Nonaka, I. (1996). *Relentless: The Japanese Way of Marketing*, New York: Harper Business.

Klir, G.J. (1976). Identification of Generative Structures in Empirical Data. *International Journal of General Systems*, 3, 89–104.

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

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

Levis, A.J. (1977). The Formal Theory of Behavior. *Journal of Social Psychiatry*, 23(2).

Maes, P. (1995). Modeling Adaptive Autonomous Agents. In C.E. Langton (Ed.), *Artificial Life: An Overview*, Cambridge, MA: MIT Press, pp. 135–162.

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.

Maturana, H.R. (1975). The Organization of the Living: A Theory of the Living Organization. *International Journal Man-Machine Studies*, 7(3), 313–332.

Maturana, H.R. (1981). Autopoiesis. In M. Zeleny (Ed.), *Autopoiesis: A Theory of Living Organization*, pp. 21–33, North-Holland: New York.

Monod, J. (1972). *Chance and Necessity*, London: Collins.

Newell, A. (1982). The Knowledge Level. *Artificial Intelligence*, 18(1), 87–127.

Nonaka, I. and Takeuchi, H. (1995). *The Knowledge-Creating Company*, Oxford: Oxford University Press.

Novak, J.D. and Gowin, D.B. (1984). *Learning How To Learn*, New York: Cambridge University Press.

Ollason, J.G. (1991). What is This Stuff Called Fitness? *Biology and Philosophy*, 6(1), 81–92.

Peirce, C.S. (1898). *Reasoning and the Logic of Things*, Cambridge, MA: Harvard University Press.

Popper, K.R. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*, London: Routledge and Kegan Paul.

Prigogine, I. (1984). *Order Out of Chaos*, Toronto: Bantam.

Rasiowa, H. and Sikorski, R. (1970). *The Mathematics of Metamathematics*, Warsaw: PWN.

Ruelle, D. (1989). *Chaotic Evolution and Strange Attractors*, Cambridge, UK: Cambridge University Press.

Sharif, N. (1978). An Introduction to Modeling. In N. Sharif and P. Adulbhan (Eds.), *Systems Modeling for Decision Making* (pp. 1–21). Bangkok: Asian Institute of Technology (distributed by Pergamon Press).

Snooks, G.D. (1996). *The Dynamic Society: Exploring the Sources of Global Change*, London: Routledge.

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

Tarski, A. (1930). Fundamentale Begriffe der Methodologie der deduktiven Wissenschaften I. *Monatahefte für Mathematik und Physik*, 37, 361–404.

Weiß, G. and Sen, S. (Ed.) (1996). *Adaption and Learning in Multi-Agent Systems*, Berlin: Springer.

Wiener, N. (1948). *Cybernetics or Control and Communication in Animal and Machine*, Cambridge, MA: MIT Press.

Wójcicki, R. (1988). *Theory of Logical Calculi*, Dordrecht: Kluwer.

Wooldridge, M., Müller, J.P., and Tambe, M. (Ed.) (1996). *Intelligent Agents II*, Berlin: Springer.

Zadeh, L.A. (1964). The Concept of State in System Theory. In M.D. Mesarovic (Ed.), *Views on General Systems Theory: Proceedings of the Second Systems Symposium at Case Institute of Technology*. New York: John Wiley.

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.