

Knowledge Level Modeling of Agents, Organizations and Technologies

Brian R. Gaines
University of Calgary
Calgary, Alberta, Canada T2N 1N4
gaines@cpsc.ucalgary.ca, <http://ksi.cpsc.ucalgary.ca/KSI/>

From: AAAI Technical Report SS-97-01. Compilation copyright © 1997, AAAI (www.aaai.org). All rights reserved.

Abstract

Knowledge management has emerged as a major industrial focus and has obvious pragmatic interpretations in terms of enterprise and workflow modeling. However, a principled approach to the management of the knowledge processes of organizations combining people and technologies requires an operational definition of *knowledge*. This article develops a knowledge level analysis of the emergence of the knowledge construct through modeling and management processes in societies of adaptive agents. 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.

1 Introduction

The objective of the research described in this paper is to derive fundamental principles for knowledge management 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 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.

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.

2 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). 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 (Dendrinis 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.

2.1 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.

2.2 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, its 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 investigate each of these requirements in turn.

3 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.

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:

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

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:

$$\forall t, u \in T, t \geq u \equiv \forall a \in A, a \circ t \rightarrow a \circ u \quad (2)$$

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:

$$\forall t, u \in T, t \geq u \equiv K(t) \supset K(u) \quad (3)$$

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

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

The development to this stage parallels that of knowledge spaces as defined by Falmagne, Koppen, Villano, Doignon and Johannesen (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:

$$x \subset K \rightarrow x \subset G(x) \quad (5)$$

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

$$x \subset K \rightarrow G(G(x)) \subset G(x) \quad (6)$$

—additional knowledge generates additional knowledge:

$$x \subset y \subset K \rightarrow G(x) \subset G(y) \quad (7)$$

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 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.

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.

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

4 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:

$$a \in b \rightarrow G(K(a)) \subset G(K(b)) \quad (8)$$

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 reevaluate 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:-

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

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

4.1 Impact of Technology on Knowledge

The measurement of the impact of organizing agents in equation (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:

$$E(a, a') \equiv G(K(a')) - G(K(a)) \quad (10)$$

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 equations (9) and (10):

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

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

5 Learning

Presupposition P1 in Section 3, 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,b).

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. For example, a related article shows that the optimization of learning to accelerate a sigmoidal learning curve to become a linear one can be managed with surprisingly weak models of the knowledge states of the agents involved (Gaines, 1996).

6 Conclusions

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 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.

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.

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.
- Borenstein, S. and Radman, Z.A. (1985). **Learning to Learn: An Approach to Study Skills**. Dubuque, Kendall/Hunt.
- Bridges, W. (1991). **Managing Transitions: Making the Most of Change**. Reading, Massachusetts, Addison-Wesley.
- Cannon, W.B. (1932). **The Wisdom of the Body**. London,
- Dendrinis, D.S. and Sonis, M. (1990). **Chaos and Socio-Spatial Dynamics**. Berlin, Springer.
- Dennett, D.C. (1987). **The Intentional Stance**. MIT Press, Cambridge, Massachusetts.
- 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. (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 & 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. Caudhill, M. and Butler, C., Ed. **Proceedings of IEEE First International Conference on Neural Networks. Vol.2.** pp.387-394.
- 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. (1996). Adaptive interactions in societies of agents. **Workshop on Intelligent Adaptive Agents.** pp.4-10. Menlo Park, CA, AAAI.
- 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.
- 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, HarperBusiness.
- Levis, A.J. (1977). The formal theory of behavior. **Journal of Social Psychiatry** 23(2)
- 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. Zeleny, M., 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.
- 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, Massachusetts, Harvard University Press.
- Popper, K.R. (1963). **Conjectures and Refutations: The Growth of Scientific Knowledge.** London, Routledge & 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. Sharif, N. and Adulbhan, P., Ed. **Systems Modeling for Decision Making.** pp.1-21. Bangkok, Asian Institute of Technology (distributed by Pergamon Press).
- Tarski, A. (1930). Fundamentale Begriffe der Methodologie der deduktiven Wissenschaften I. **Monatshefte für Mathematik und Physik** 37 361-404.
- Wiener, N. (1948). **Cybernetics or Control and Communication in Animal and Machine.** Cambridge, Massachusetts, MIT Press.
- Wójcicki, R. (1988). **Theory of Logical Calculi.** Dordrecht, Kluwer.
- Zadeh, L.A. (1964). The concept of state in system theory. Mesarovic, M.D., Ed. **Views on General Systems Theory: Proceedings of the Second Systems Symposium at Case Institute of Technology.** New York, John Wiley.