WHEN DOES A BIZARRE SYSTEM HAVE THE FEELING OF KNOWING?

STEPHEN W. KERCEL
Oak Ridge National Laboratory
P. O. Box 2008 MS 6011
Oak Ridge TN 38731-6011

ABSTRACT

Can an instrument tell whether a process "feels right?" No such instrument has ever been constructed. Nevertheless, there is no fundamental reason to suppose that it cannot be done. It is expected that cognition will be one of the breakthroughs in computer science in the near future.

INTRODUCTION: PATTERN RECOGNITION VS. COGNITION

It is expected that man-made cognitive systems will be one of the major technological breakthroughs in the next 10 to 20 years. Indeed, instrumentation systems based on cognitive principles would lead to a revolution in sensing and control. Such an instrument would tell whether or not a process "feels right." This is a capability that would dramatically exceed the conventionally accepted fundamental limits on instrumentation.

Imagine the application of cognitive systems to the monitoring of manufacturing processes. They would enable the implementation of anticipatory maintenance, allowing defective components in production systems to be replaced just before they fail. Their sense of "not feeling right" would anticipate catastrophic occurrences, such as breakouts in continuous casting, and provide sufficient warning to enable action to avert the catastrophe. They would be able to make on-line real-time observations of physical states not currently accessible, such as the time-evolution of interacting phases in solidifying molten metal.

Making conventional measurements with conventional instruments is fundamentally limited, because these instruments are intended for strictly reductionistic (occurrences projected onto a list of lists of numbers) observations. In fact, we go to great cost and effort to make instruments (and the processes that they monitor) behave in a reductionistic manner. Thus, a fundamental limit on conventional measurement is that any entailment in any natural system that cannot be reduced to vectors of numbers will go completely unnoticed and unreported.

The fundamental problem is the generic problem of pattern recognition, or classification (Rosen, 2000, pp. 136-140, Bateson, 2000). Is the exemplar, x, a member of the class, X? To qualify as a member of the class X, the exemplar, x, must possess the distinguishing features of class, X. But, what are the distinguishing features? As both a practical and philosophical matter, *how do we know* whether or not x has them?

The logical question is "Is P(x) true?" where P(x) is the proposition $x \in X$. We try to answer this question by building a feature detector, Q, that recognizes P, and determines whether or not P(x) is true. However, the reading, Q(y), produced by the feature detector is yet another proposition, and we must validate Q(y) before we can say $Q(y) \rightarrow P(x)$. Suppose we do this with a feature-detector detector, R, in the hope that $R(z) \rightarrow Q(y) \rightarrow P(x)$. Unfortunately, we always need one more detector to validate the next one back in the progression; total validation would require an infinite regress.

One way out of this predicament is to have a model of y constructed from independent information. From the model we could test the truth of Q(y), and evade the infinite regress. However, suppose there is no independent source of information about y? If it requires an infinite regress to validate P(x), it would still require an infinite regress to validate Q(y). Short of mystical revelation, we will encounter this problem every time we try to construct any system (abstract model or physical instrument) to measure any other system.

The only other way out of the dilemma is to make the regress finite by folding it back on itself. We can limit ourselves to two stages of regression by forcing z = x. Then we could say $R(x) \rightarrow Q(y) \rightarrow P(x)$. However, under this constraint of self-reference, propositions about x and y cannot be separated. We have formed an impredicative loop. As is widely recognized in computer science, the cost of using an impredicative loop to break the infinite regress is that no algorithm can be devised to implement it. As in Rosen's equation 8.6, $x \leftrightarrow y$, or x maps to y maps to x. There exist incomputable models of x and y with semantic meanings not captured by syntax (a list of symbols and manipulation rules).

Faced with the apparent Hobson's choice between a predicative infinite regress and an impredicative (and incomputable) loop, conventional pattern recognition opts for the predicative regress. Immediately recognizing that a real-world computer cannot count to infinity, this regress has to be truncated, typically right after the first step. In other words, R(z) or Q(y) are validated by *assuming* them to be true, whether explicitly or (more often) implicitly. However, we typically have no warning as to when the assumption will break down. The assumption of validity in the absence of supporting observations is one reason for the notorious brittleness of neural nets. This brittleness is reflected in the fact that a classifier will work with a very large data set and make very few errors, but then fail miserably when confronted with a new, but seemingly similar data set. This type of failure can be reproduced at will; it happens consistently whenever the system is first demonstrated to the customer.

What about the other alternative, the self-referential loop? As will be seen in subsequent sections of this paper, impredicatives are at the foundation of cognition. This is the way that nature has solved the classification problem in living minds. But aren't impredicative loops irreducible and incomputable? They are, but that does not mean that they are intractable. There exists a whole realm of irreductionist mathematics that is no less logical than the reductionist branch of mathematics familiar to engineers. Self-referential behavior such as cognition, being irreducible, is bizarre. However, it is logically tractable, and therefore not absurd.

In exploring the concepts of bizarre systems, two explicit but unprovable assumptions are made at the foundation. One is causality, the notion that events in reality are caused by other events. The other is logic, the notion that propositions in a formal system are implied by other propositions.

CONGRUENCY OF FORMAL AND NATURAL BIZARRE SYSTEMS

Faced with two unattractive alternatives, infinite regress or impredicative loop, we ask, "How did nature solve the problem?" A considerable body of experimental evidence in both neurology and psychology suggests that nature's solution to the classification problem is model-based (Freeman, 2000, 1999, Damasio, 1999, Dilts, et al., 1980, Caulfield et al., 2000, 1999).

But then, what is a model? In describing the properties of the Modeling Relation, Rosen tells us that modeling "is the art of bringing entailment structures into congruence" (Rosen, 1991, p. 152). However this statement leaves us wondering what he really means. How does art enter into the picture; are we not instead supposed to be scientific? What is an entailment, much less an entailment structure? What does it mean that two different entailment structures are congruent? In what sense are they not identical? If they are not identical, what similarity between them causes us to declare them congruent? Why do these questions matter? If they do matter, then what should we do about them?

The first point to appreciate is that the Modeling Relation is a relation in the formal mathematical sense (Rosen, 1999, p. 374). Suppose that A and B are sets, and that there exists a set, R, of ordered pairs, where the first element of each pair in R is an element of A, and the second element of each pair in R is an element of B. In mathematical notation: $a \in A$, $b \in B$, $(a,b) \in R \Leftrightarrow aRb$. In Rosen's Modeling Relation, the members a and b of each ordered pair in R are entailments from two different systems.

Entailments are the consequences of the order or organization of a system. There are two sorts of systems that might provide entailments to the Modeling Relation, natural systems and formal systems. Natural systems are systems in physical reality that have causal linkages; if certain causative events impinge upon a natural system, then the system will behave in a certain way, or produce certain events in effect. This consequential linkage of cause and effect in a natural system is a causal entailment. Formal systems are conceptual systems that have inferential linkages; if certain hypothetical propositions impinge upon a formal system, then they will produce certain consequential propositions in conclusion. This consequential linkage of hypothesis and conclusion in a formal system is an inferential entailment. Entailment structures are inherent within a system; they are the distinguishing features that characterize the system (Rosen, 1991, p. 98). They do not cross over from one system to another.

This is represented in Fig. 1, where we see a natural system, N, distinguished by its structure of causal entailments, a, and a formal system, F, distinguished by its structure of inferential entailments, b. The entailment structures of two distinct systems are distinct from one another; causes or hypotheses in one do not produce effects or conclusions in the other. In fact, this provides the answer to one of the questions posed above. Its self-contained entailment structure is what provides identity to a system and distinguishes it from other systems. This is the urgently sought-after distinguishing feature of the classification problem.

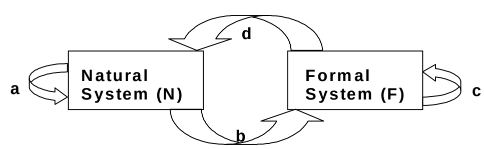

MR={(a,c) | a = b + c + d}

Figure 1. The Modeling Relation

The fact that distinct systems are non-identical does not preclude them from being regarded as being in some sense similar. Similar systems should have distinguishing features that closely correspond to each other. Dissimilar systems should have distinguishing features that do not closely correspond to each other. As already noted, the distinguishing feature of a system is its entailment structure. Thus, we would expect similar systems to have entailment structures in which there is some degree of correspondence between the entailments.

To establish this correspondence, consider a system of encodings and decodings (Rosen, 1991, p. 59). For example, we might have a system of encodings that encodes a set of events in the natural system, N, in Fig. 1, into a set of propositions in the formal system,

F. We might also have a system of decodings that decodes a set of propositions in the formal system, F, into a set of phenomena in the natural system, N. Although the two systems remain independent in the sense that causes or hypotheses in one do not produce effects or conclusions in the other, the two systems can be linked by encodings and decodings.

This linkage between entailment structures provides the means of determining the similarity between two systems. Suppose that an event, e_1, in N can be encoded to a proposition, p_1, in F; we can think of the encoding arrow, b, in Fig. 1 as a measurement on a natural system. Suppose further that the proposition, p_1, when applied as a hypothesis in the inferential structure in F entails another proposition, p_2, in F as a conclusion. In other words, the two propositions are entailed as an implication, $c = (p_1 \rightarrow p_2)$, in F. Suppose that this entailed proposition, p_2, in F can be decoded into an event, e_2, in N; we can think of the decoding arrow, d, in Fig. 1 as a prediction by a formal system.

Rosen defines congruency between the entailment structures in the following way (Rosen, 1991, p. 61). Suppose that in the underlying reality, the event e_1 in N causes event e_2 in N. In other words, the two events are entailed as a causal linkage, $a = (e_1 \rightarrow e_2)$, in N. Suppose further that the linkages commute. Event e_1 is encoded by b to proposition p_1, which implies proposition p_2, which decodes to event e_2, and that there is exact correspondence between the predicted event e_2, and the caused event e_2. The commutation is also described as $a = b+c+d$. (Note: In this context, + is the symbol for concatenation.) If there exists no such entailment c in F, having a commutative relationship with some entailment a in N, then the two systems do not have congruent entailment structures. Entailment structures are congruent to the extent that such correspondences between entailments exist. If such correspondences between the entailments in the two systems do exist, then we can learn something about one entailment structure by observing the other. This is the essence of Rosen's Modeling Relation. When it is applied to a formal system to obtain predictions about a natural system, the inferential entailments in the formal system correspond to the causal entailments in the natural system. Where the relationship holds up and where it breaks down are both understood. This is where the Modeling Relation differs from "black box" approaches. Construction of a "black box" makes no claims about the causal links in underlying reality, offers *no understanding* of the natural system it purports to describe, and offers no warning as to when the description will break down.

In contrast, for any valid Modeling Relation, the identification of the encodings and decodings between two systems is an act of discovery based on insight or understanding. The benefit of this understanding is the awareness of the specific entailments so described, and a clear indication of the scope of applicability (or non-applicability) of the formal system as a model. The cost of this understanding is that it is an art and not a science in the reductionist sense; there is no automatic or algorithmic method for determining either the encodings or decodings. In fact, there is not even any necessity or assurance that the system of decodings can be obtained from some straightforward inversion of the encodings.

Both the cost and benefit of Rosen's formalism go to the question of why it matters. We attempt to use formal systems to learn about natural systems because simply observing the natural system is too slow, too costly, too dangerous, or too inaccessible. Formalism without understanding of the underlying reality is seductive. At a fraction of the cost of the real thing, it offers the illusion of understanding. However, pretending that a modeling relation exists between a formal system and a natural system when there is no congruency of entailment (sometimes called an "as if" model) gives us results that we cannot trust, and on which we cannot afford to risk lives, safety, or vast sums of money.

Rosen draws a distinction between modeling and simulation. Modeling involves the discovery of congruent entailments. Simulacra, or "black box" algorithms, find coincidences in observations and proceed from the hope that they are extrapolatable. The distinction is important because in real-world settings models are far more trustworthy than simulations.

In describing the models constructed by consciousness, what Rosen calls a model, Dilts calls an "elegant model," one whose elements, or entailments are causally important (Dilts, *et al*., 1980).

What does the modeling relation tell us about instrumentation? Suppose the left-hand box in Fig. 1 is the abstraction of some natural system. Further suppose that the entailment structure, a, includes some causal entailment of the natural system that we would like to observe. Then the instrument is abstracted by the right-hand box, and the entailment structure, c, includes some entailment of the instrument that is congruent with the causal entailment that we would like to observe. To the present day, for every instrument ever built, the only entailments actually used for measurements are those that are also congruent with some entailment within a vector space of real numbers.

However, there are several facts that are seldom called to the attention of instrumentation engineers. First is the fact that there exist non-denumerable entailments in natural systems. Second is the fact that there are formal systems that have inferential entailments that are irreducible to vectors of numbers, but are nevertheless congruent with causal entailments and are amenable to logical manipulation. Because these things exist, it is possible to make predictions about non-denumerable processes in natural systems without recourse to mysticism.

There remains the inconvenient technological issue that irreducible mathematical constructs cannot be run on a conventional computer using algorithmic programs; irreducible entailments are not Turing computable. However, there are two possible ways to evade this limitation. One, suggested by Landauer, is not to assume that Turing had the last word on computing; perhaps an irreducible model of the digital computer, more universal than the Universal Turing Machine can be discovered (Landauer and Bellman, 1999). Alternatively, a number of researchers are investigating physical effects that include causal entailments that are irreducible (Dress, 1999). These effects may lead to a different type of functional component that could be used in the right-hand box of the modeling relation to manipulate entailments that are congruent to the irreducible entailments in natural systems that the present generation of instruments completely fails to capture. Man-made functional components do not presently exist, but they may be coming soon (in years rather than decades).

THE PHYSIOLOGICAL FEELING OF KNOWING

Whether physicists eventually cobble up a quantum- entangled functional component, or biologists concoct an artificial wetware structure that behaves something like a brain, the substrate is necessary but not sufficient for cognition. Refining our question about how nature does it, what we really want to know is how nature organizes the activity of such a functional component to produce cognition.

From the perspective of neurophysiology, cognitive behavior is irreducible to a list of lists of numbers. This being the case, no list of lists of numbers, no matter how big, can model intelligent behavior. To discuss cognition on any level deeper than merely tabulating empirical observations, some logical description of it must be found that is not limited by the mathematical laws of vector spaces.

The evidence of neurologists such as Walter Freeman and Antonio Damasio suggests that the mind is more than a direct sum of a set of parts, or has a nature along the lines of the Gestalt (Freeman, 2000, 1999, Damasio, 1999, p. 88). A formal model (as in Rosen's Modeling Relation) of the brain/mind natural system based on the claim that the mind is a direct sum of a set of parts would ignore the key features of the operation of the mind. In other words, consciousness is irreducible.

Biological consciousness is a million times slower than computer operations. Neural firing rates are on the order of milliseconds; recently available desktop computers perform operations on the order of nanoseconds. This extreme slowness has two practical consequences. First, the animal avoids conscious direction of action wherever possible. Second, the only way for such a slow system to provide useful direction is to be anticipatory (Caulfield, *et al*., 1999). In other words, consciousness is driven by final cause (Dilts, 1994, p.26). Reductionism ignores (or more often militantly denies) the existence of final cause.

Rather than the old cliche "Seeing is believing" the process actually works the other way around, "Believing is seeing." The brain sees what it believes it sees. In fact, if the brain cannot construct a meaningful model from sensory input, it often ignores it (Caulfield, *et al*., 1999). For example, victims of blindsight are quite insistent that they cannot see an object, but can correctly point to its location, *and have no notion as to why they know* (Damasio, 1999, p. 268).

What do brains do that other functional components do not? They provide to their owner an answer to the question, "what do I do next?" In particular, consciousness provides its owner with a useful approach to unforeseen circumstances (Caulfield, *et al*., 1999). The limbic system ultimately conveys this to our awareness as how we feel about a given situation. Similarly, an engineered conscious artifact would evaluate its situation in terms of an overall feeling.

Intentionality is active, and seeks guidance for future action (Freeman, 1999). In particular it has three distinguishing features, intent, wholeness, and unity. It differs from classical physical systems in all three particulars. Intent produces effects from final cause; classical systems ignore final cause. A whole system is irreducible to parts; a classical system is the direct sum of its parts. A unified system is context dependent, and retains its identity even though it changes in response to changes in the ambience; a classical system is context independent. By the classical concepts of science, intentionality is wholly unscientific. Do we ignore intentionality because it is "unscientific," or do we expand the scope of science to examine phenomena previously ignored?

Freeman shows that the limbic system is the organ of intentionality (Freeman, 1999). A macroscopic brain state is formed by a system of interacting neurons. The limbic system uses these brain states to generate self-organized spatio-temporal patterns. These lead to goal formation and the direction of behavior to meet goals. Awareness results from an ongoing cycle of preafference (preparing sensory cortices for the consequences of motor actions), perception (abstraction of sensations), and update (integration in the hippocampus of preafference and percepts into the mental state).

From the perspective of neurophysiology, the state of a brain is the description of what it is doing in some specified time period (Freeman, 1999). From the perspective of psychology, state is "the total ongoing mental and physical conditions from which a person is acting" (Dilts, 1994, p. 314). It is noteworthy that these descriptions of state are very similar, and are not restricted to the concept of the "state vector" of classical physics. Short-term brain states as characterized by patterns of populations of neurons are irregular both in space and time. They resemble a hurricane more than a march of symbols in a Turing machine.

As Freeman notes, "Brain systems operate on many levels of organization, each with its own scales of time and space. Dynamics is applicable to every level, from the atomic to the molecular, and from macromolecular organelles to the neurons into which they are incorporated. In turn the neurons form populations and systems, and so on up to embodied brains interacting purposively with their environments. Each level is "macroscopic" to the one below it and "microscopic" to the one above it. Among the most difficult tasks are those of conceiving and describing the exchanges between levels, seeing that the measures of time and distance are incommensurate, and that causal inference is far more ambiguous between than within levels" (Freeman, 2000).

But, how do we exchange meaning between levels? Freeman says that we must "conceive, identify and model an intervening 'mesoscopic' level, which is the self-organizing neural population." However, this model has some rather peculiar properties. As Freeman says, "the main contribution of mesoscopic neurodynamics is to the self-organizing chaotic patterns of gamma activity in primary sensory areas and the limbic system, which make perception a creative act of neural masses." However, this chaos through which meaning is exchanged is much more complex than the "deterministic chaos" of reducible non-linear systems driven at chaotic operating points. It is a self-referential chaos, and hence, irreducible and non-Turing computable.

It is fairly widely recognized that organisms with brains relate to one another through their limbic systems (Caulfield, 2000). The relationship might be hierarchical or operate among entities on the same level. Caulfield argues that in control problems for distributed systems (as might be used in conventionally engineered systems), the hierarchical model is more appropriate. However, this is not how natural intelligence operates. Instead "they pass summaries of the world *and themselves* (emphasis added) to the next level. Those summaries are the world of the next level. In the simplest hierarchical system, only the bottom layer interacts with the real world. Same-level interaction can be arranged in the same way. Global control comes from the top layer in terms of global concerns. All other levels are involved in implementation."

The bottom layer is directly stimulated by physical causes, and in resulting effect, produces a stream of chemical and electrical markers. Through various levels of the hierarchy, the data in the markers are abstracted, until at the highest level, the abstraction conveys meaning. What is it that confirms this meaning to the mind doing the abstracting? The confirmation comes when it feels right, emotively (Damasio, 1999, pp. 312-315).

THE PSYCHOLOGICAL FEELING OF KNOWING

Perception is based on interacting models of ourselves and of our ambience (Caulfield, *et al*., 1999). The patterns of sensory input assimilated at the unconscious level of the biological mind can be externally observed and expressed concisely (but not reducibly) in representational models, in three interacting basic representation systems, auditory, visual and kinesthetic (Brown-VanHoozer, et al, 2000).

In functioning human minds, sensory information is abstracted into percepts via an elegant automatic process at the unconscious level. We can, by a willed choice, consciously abstract these percepts into concepts, and use them to manipulate learning, memory, language, knowledge belief systems and willed behavior (Brown-VanHoozer, 2000). The strategies and patterns assimilated at the unconscious level can be externally observed and expressed concisely as models for application.

At this level of abstraction, we transform content into processes that can be expressed in terms of the three basic representation systems and the entangled interactions between them. These are the foundation for how effective choices and belief systems are generated through sensory derived processes. They provide the means by which learning strategies are constructed, how memory is accessed, stored, retrieved or recalled, and how knowledge is actualized at the conscious and unconscious level.

This actualization of knowledge is accomplished through a feedback process characterized as test-operate-test-exit (TOTE) as originally formulated by Miller, Galanter, and Pribram (Dilts, 1994, pp. 290-291). The "operate" consists of applying the sensation(s) to various combinations of the representational models. The "test" is the question of whether or not the knowledge feels right, and updating the models if it feels wrong. The kinesthetic-emotive representation system performs the test, and it keeps feeding back into the TOTE process until it either feels right or quits (Brown-VanHoozer, 1999). To "exit" the TOTE loop means to use the newly abstracted knowledge to proceed toward some internally defined goal or final cause (Dilts, 1994, pp. 41-44). We perceive the result of the final test of the TOTE process as a comforting and satisfying feeling of knowing (Damasio, 1999).

Disruption of these representation systems causes terrible emotional distress. For example, chronically blind people whose sight is suddenly restored by medical intervention typically find themselves overwhelmed by a jumble of sensation that is meaningless to the auditory/kinesthetic representation system that their mind has constructed to survive in a sightless world. The distress is so great that few survive the catastrophe of acquiring sight in adulthood for more than a few years (Sacks, 1995, pp. 142-151).

In autism, the sensory wetware is present and produces the right markers, but due to a neurological malfunction, the representation systems do not construct sensible percepts from them; as a result the typical autistic child lives a life of unremitting terror (Sacks, 1995, pp. 253-255). Autistic adults often exclusively "think in pictures," and verbal language is a pure abstraction into which high-functioning autistics have learned to translate their visual concepts. They have enormous difficulty imagining that other people do not think in pictures, and find it difficult or impossible to relate to people who do not. Their highly developed visual representation systems overstimulate the neurological "fight or flight" response, and the other representation systems are easily overwhelmed by sounds, touches, tastes and smells. As a result, the daily life of the autistic adult is only slightly less frightening than that of the autistic child (Sacks, 1995, pp. 269-271). It is not surprising that if the abstraction of sensations into "knowing" feels right, then a meaningless jumble of sensations feels terribly wrong.

CONCLUSION: THE CONSTRUCTION OF IRREDUCIBLE KNOWING

There are seven important issues in the implementation of intelligent systems. First, how do you implement self-reference? This involves developing either a computing component that includes self-referential behavior, or a self-referential model of computation that does not suffer from the limitations of the Turing machine. Second, how do you extend a mathematical structure to incorporate information about its context? Third, how do you move mathematical structures to a new context and assess their validity? Fourth, how do you define structures before you define the elements of the structures? Fifth, how do you capture the modeling process in a mathematical structure? Sixth, how do you decide when your notational system is inadequate? Seventh, how do you fix the inadequacy? It does not appear that *any* of these seven questions can be addressed within the bounds of reductionist mathematics. Remarkably, this does not appear to be a fundamental problem; all these questions appear to be tractable with irreducible mathematics (Landauer and Bellman, 1999).

Intelligent behavior is bizarre, requiring among other things that the intelligence be able to recognize whether or not its own self-referred sense of a process "feels right." However, it is not absurd. Absurd behavior is unconstrained by any laws of natural systems and not amenable to description by a logical formal system. In contrast, bizarre behavior is merely counterintuitive. It remains congruent with causality, and can be modeled with logically tractable, albeit impredicative mathematics. The bizarreness is reflected in the fact that intelligence cannot be subsumed in an algorithmic program. Given a substrate that has the potential to support consciousness, the task of artificially imposing that consciousness upon it is a something new in the art of engineering. Nevertheless engineers do have guidance as to how to set about it. The mathematical description of the self-referential processes in naturally occurring wetware should provide the necessary insight into how a practical man-made mind could be made and operated.

ACKNOWLEDGMENTS

This research was performed at Oak Ridge National Laboratory, which is operated by UT-Battelle, LLC, for the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. I thank Mr. Roger Kisner, Dr. Glenn Allgood, Dr. W.R. VanHoozer, and Dr. Alenka Brown-VanHoozer for their very helpful comments.

REFERENCES

Bateson, G., Steps to an ecology of the mind: Collected essays in anthropology, psychiatry, evolution, and epistemology, University of Chicago Press, 2000.

Brown-VanHoozer, S.A., 1999, "Models of Reality," in ANNIE 99, pp. 199-205.

Brown-VanHoozer, S.A., and VanHoozer, W.R., 2000, "Models of Reality: The Unconscious Tells All," this volume.

Brown-VanHoozer, S.A., Kercel, S.W., and VanHoozer, W.R., 2000, "The Model-Based Mind," in Proceedings of SMC 2000: IEEE International Conference on Systems, Man, and Cybernetics.

Caulfield, H.J., Johnson, J.L., Schamshula, M., Inguva, R., and Egarievwe, S.U., 1999, "Perception, Attention, and Consciousness in Human Artifacts," in ANNIE 99, pp. 237-246.

Caulfield, H. J., 2000, "Beyond the Anthill: Consciousness in Networks of Conscious Entities," this volume.

Damasio, A .R., 1999, *The Feeling of What Happens*, Harcourt, Brace and Company, pp. 67-71.

Dilts, R.B., 1994, *Strategies of Genius*, Meta Publications, Capitola.

Dilts, R.B., Grinder, J., Bandler, R., and DeLozier, J., 1980, *Neuro-Linguistic Programming: Volume I The Study of the Structure of Subjective Experience*, Meta Publications, Capitola.

Dress, W.B., 1999, "Epistemology and the Rosen Modeling Relation," in *ANNIE 99*, pp. 231-236.

Freeman, W. J., 1999, "The Neurodynamics of Intentionality is the Basis of Intelligent Behavior," in *ANNIE 99*, pp. 185-191.

Freeman, W. J., 2000, "Intention and Consciousness Emerge through Mesoscopic Brain Dynamics," this volume.

Landauer, C., and Bellman, K., "Architectures for Embodied Intelligence," in *ANNIE 99*, pp. 215-220.

Rosen, R., 2000, *Essays on life itself*, Columbia University Press, New York.

Rosen, K., 1999, *Discrete mathematics and its applications*, Fourth Edition, WCB/McGraw-Hill.

Rosen, R., 1991, *Life itself: a comprehensive inquiry into the nature, origin and fabrication of life*, Columbia University Press, New York.

Sacks, O., 1995, *An anthropologist on Mars*, Vintage Books, New York.