Behavioral Personal Digital Assistants: The Seventh Generation of Computing

Kenneth R. Stephens and William R. Hutchison
BehavHeuristics, Inc., College Park, MD

Skinner (1985) described two divergent approaches to developing computer systems that would behave with some approximation to intelligence. The first approach, which corresponds to the mainstream of artificial intelligence and expert systems, models intelligence as a set of production rules that incorporate knowledge and a set of heuristics for inference and symbol manipulation. The alternative is a system that models the behavioral repertoire as a network of associations between antecedent stimuli and operants, and adapts when supplied with reinforcement. The latter approach is consistent with developments in the field of "neural networks." The authors describe how an existing adaptive network software system, based on behavior analysis and developed since 1983, can be extended to provide a new generation of software systems capable of acquiring verbal behavior. This effort will require the collaboration of the academic and commercial sectors of the behavioral community, but the end result will enable a generational change in computer systems and support for behavior analytic concepts.

The progress of computer science has often been divided into a number of phases or generations. The first generation systems were simple machines that were programmed in the most fundamental machine language. The second generation made programming much more practical by introducing assemblers and compilers. The third generation added databases. The fourth generation made the databases more productive and usable. "Fifth generation" is synonymous with symbolic artificial intelligence. The term "sixth generation" has been used to refer to neural networks. We propose that the next major breakthrough will be in intelligent systems that exhibit verbal behavior, therefore we refer to the effort to produce such systems as the Seventh Generation (7G) project.

TWO APPROACHES TO INTELLIGENT SOFTWARE

For over three decades, computer scien-

tists and others have pursued the goal of developing "intelligent" machines. Machines that could "think instead of merely compute" would be useful for any number of tasks which are dangerous to humans, or where intelligent automation can cut costs or generate additional revenue, or in which human expertise is rare and should be captured and replicated, or in which the efficiency and speed of a computer are desirable but where the end result should be similar to that which a human would produce. Skinner (1985) pointed out that there are two approaches to developing systems which would be considered intelligent: a cognitive approach based on theories of symbol manipulation and representation of knowledge as rules, and a more dynamic approach based on the emergence of contingency-shaped behavior.

The goal of building an intelligent system has been difficult to achieve, partly because it is problematic to define intelligence in any context, as Schlinger (1992) has pointed out in this volume. From its beginning, the field of artificial intelligence (AI) has variously defined intelligence to

Reprint requests may be addressed to either author at BehavHeuristics, Inc., 335 Paint Branch Drive, College Park, MD 20742. The authors gratefully acknowledge the editorial assistance of Ken Otwell, Bruce Lombardi, Sharon Hornby, Kurt Rottweiler, Paul Werbos, and Obbie Pet.

include the abilities to reason, to solve problems by utilizing heuristics (rules-of-thumb) rather than algorithms, to represent knowledge, to manipulate symbols, or in general, to think (and therefore behave) in a manner similar to humans. The theoretical guidance for this field in recent years has come from cognitive science (a discipline formed by the marriage of cognitive psychology and computer science). Based on public perception and the attitudes of the AI community, the objective of developing intelligent machines would seem to belong to this group by default. This is not entirely true. There has been noteworthy progress by others, including those in cybernetics and behavior analysis.¹

One criterion for intelligence that is often omitted by cognitive scientists is the ability to learn and adapt. This reveals a fundamental paradigm difference between cognitive science and behavior analysis. To the cognitive scientists, learning is the acquisition of rules and knowledge, which can then be stored in a "rule base" or other knowledge structures of the intelligent system. Most of the work in machine learning, therefore, has been focused on induction of rules from data. In practice, the rules of most "expert systems" are simply supplied by a "knowledge engineer." Without the knowledge engineer, most expert systems would be lacking an essential component of an intelligent system.

There have been well-publicized successes of rule-based expert systems in specific problem domains (like certain types of medical diagnosis or configuration of large computers) which are especially suited to a symbolic, logical approach. But there has also been disappointment, as symbolic AI has often failed to live up to the excessive promises which were made for it in the early 1980s.

To a behavior analyst interested in intelligent software, learning is the development of differential stimulus control over the behavior of the system. Except for the

use of "certainty factors" which give some flexibility to the relationship between conditions and actions, or between premises and conclusions, the production rules in expert systems represent rigid relationships. These rules of the cognitive scientists are, ironically, more like the "stimulus-response" psychology they criticize as "behaviorist" than to the three-term contingency of radical behaviorism. Contingency-shaped behavior is often complex and probabilistic, and stimulus control is usually "influential" rather than "reflexive." However, rules by themselves are insufficient to *describe* dynamic behavior, and therefore are inadequate as the basis for learning and adaptive behavior. Many of the most interesting challenges for intelligent systems come in domains where control is complex or probabilistic or constantly changing. An adaptive system must be able to come under emerging (but partial) control, or competitive control from different sources. Further, in complex situations, behavior is often under complex stimulus control from multiple sources.

Beyond the issue of whether rules can *describe* behavior, there is a need to understand the process by which stated rules may influence behavior. That process has not been addressed at all in AI, to our knowledge; rules are simply programmed, and they exert total control. For behavior analysts, rule governance is an important topic which has been the subject of much research and debate. This process of *learning* rules is extraordinarily important and understudied, since arguably *most* behavior of humans is rule-governed in origin.

Contingency-shaping may eventually lead to the emergence of a rule and rule governance. Conversely a rule may be a good starting point to enable the adaptive system to come into contact with contingencies, although under some conditions following a rule or instructions may reduce responsiveness to contingencies (Catania, Shimoff, & Matthews, 1989). In a critical treatment of expert systems, Dreyfus and Dreyfus (1986) point out that human learners "do not appear to leap suddenly from

¹As we have pointed out elsewhere (Stephens & Hutchison, unpublished manuscript, 1989), the latter two groups have more in common than most of their members think.

rule-guided 'knowing that' to experience-based know-how."

Over the last several decades, a type of computer system has emerged which is capable of adapting when supplied with feedback. These systems have come to be known as "neural networks," even though the community of researchers and software engineers who develop and study them is multi-disciplinary, and the neural metaphor for adaptive behavior only applies to a fraction of the implementations. While we prefer the term "adaptive network systems," we often use the term which has received the more widespread recognition. Neural networks work well in situations where "knowledge" cannot be easily stated as rules, where it is desirable to have the computer software infer the rule, where there is incomplete stimulus control (because of missing or unreliable data), where stimulus control is complex and subtle, or where relationships keep changing. While the field of neural networks actually predates symbolic AI, certain milestone events occurred in the late 1960s which shifted the emphasis to the type of system that is now usually thought of as AI. For many years, research on neural networks continued in relative obscurity and with minimal funding. In the 1980s, however, there was a dramatic resurgence of interest in neural networks. Developments of capabilities such as back propagation of error (Werbos, 1974), temporal differences (Sutton, 1988), and the "3-net" or adaptive critic model (Werbos, 1977) have enabled a new generation of commercially viable adaptive network systems.

The field of neural networks is multi-disciplinary, and there is a faction whose point of view is most similar to the cognitive science orientation within AI. These researchers often refer to their work as "connectionist AI" and they are more likely to talk about their work at AI conferences and in AI journals than are many of their colleagues. For instance, the "PDP Research Group"² (Rumelhart, McClelland, et al, 1986) generally falls within this paradigm.

As recent reviews in JEAB (cf. Donahoe

& Palmer, 1989; Kehoe, 1989), and our own papers (Hutchison, 1984, 1985a, 1985b; Hutchison & Stephens, 1986; Stephens & Hutchison, 1989); Stephens, Hutchison, & Bell, 1988), have pointed out, the neural network paradigm should be appealing to behavior analysts. Conversely, behavior analysts have much to offer to the field of neural networks. Early papers by Klopff (1982) and Sutton and Barto (1981) discussed the learning properties of neural networks in terms of reinforcement, mainly respondent conditioning. Even more interesting, however, is a type of adaptive system which explicitly incorporates principles from the experimental analysis of behavior. Such software has existed since 1983. Originally called the Operant Adaptive System™, it has now been incorporated in the BANKET™ system. The first commercial application of BANKET was an airline yield management system called Airline Marketing Assistant (Hutchison & Stephens, 1987b) which is often cited as a conspicuous success story in articles about the commercial potential of neural networks.

INTELLIGENT SYSTEMS AND VERBAL BEHAVIOR

AI researchers have also been interested in "natural language" processing (NLP), in which the goal is to apply linguistic theory so that systems are able to communicate with humans more easily than allowed by the current generation of computer operating systems (see Cherpas, 1992 for an excellent review of NLP). When the connectionist AI community has attempted to apply neural networks to this problem, it has almost always had the objective of relating human input to linguistic structures. Networks learn from a set of sentences or phrases entered as input. The role of the network is to assign a "linguistic role" and meaning to each word in the linguistic structure in such a way that logical inference can be made. Another line of research presents words for which the network must learn to output regular and irregular forms of plural nouns, third person singular verb, and so on (e.g.,

²PDP refers to parallel distributed processing.

Gupta & MacWhinney, 1992; Rumelhart & McClelland, 1986). However, there is evidence that future work by connectionist cognitive researchers may become easier to interpret in light of our analysis of verbal behavior. Hinton (1991) has begun drawing a distinction between intuitive and rational behavior, which is similar to the distinction between contingency-shaped and rule-governed behavior. St. John (1992) proposed a rather radical perspective: to train "language in service of a task." He argues that taking a functional perspective avoids several acknowledged problems of structuralist approaches (stipulating a linguistic representation as the goal of comprehension; defining the meaning of concepts; combining lexical, syntactic, and task-context constraints; and the informality of natural speech). In spite of the typical connectionist denial that they are becoming behaviorist, a paradigm shift seems to be occurring without their recognition that these positions and issues have been argued years before.

A well-known criterion for a truly intelligent computer system was first formulated many years ago by the famous computer scientist, Alan Turing: if a human observer, physically isolated and interacting with computer software through a terminal, cannot tell whether or not the other participant in a dialog is human or computer, then artificial intelligence has effectively been demonstrated. An annual competition based on the Turing Test, the Loebner Prize, has been hosted by the Cambridge Center for Behavioral Studies since 1991. While there have been significant advances in NLP, no system yet has passed the Turing test and won the Loebner Prize. As difficult as this goal is, an even more demanding objective is required to support the next generation of "personal digital assistants" or PDAs. PDAs will be software agents that will provide more effective functionality, and a much more supportive interface to the computer system. Several companies (Apple, AT&T, Phillips, SONY, Matsushita, Motorola, etc.) have already begun focusing their R&D efforts on producing these software agents, and on a

new generation of "personal communicators." PDAs will need to be able to adapt to the requirements and style of an individual system user: to present information to the user in the manner which has proven to be most effective in the past, to accept inputs in short-hand or slang, to incorporate explicit or implicit user preferences, to anticipate likely activities and then provide support for these anticipated actions, to do what the user means rather than says, and to perform some activities automatically.

This need for flexibility and responsiveness implies to a behavior analyst that the behavior of these synthetic agents will have to be acquired in much the same way as humans. In the past, we have occasionally encountered skepticism that computer input/output can ever constitute verbal behavior, or that the behavior of a human computer user should be considered verbal behavior. At the point that PDAs can functionally emulate verbal operants, this argument becomes moot. When the functional basis of the behavior of the system is sufficiently similar to that of humans, what we know about human verbal behavior and its acquisition and control becomes quite relevant in the design and development of such systems.

From our perspective as behavioral software developers, we believe it is now feasible for an adaptive software system to emulate the functional properties of verbal behavior, as described by Skinner (1957). What is proposed here is a type of PDA that incorporates behavioral principles. It will be based on the current generation of BANKET software, with overall architectural enhancements developed as needed to handle the verbal operants. We will refer to these new systems as Behavioral PDAs or BPDAs. We plan to develop BPDAs in a collaborative project between the commercial and academic segments of the behavioral community.

IDEAL BEHAVIORAL PDAS

The BPDA will have several functions. First will be to allow the end user of a computer system to interact with it in a much more relaxed fashion than current systems,

including those with NLP capabilities. Adaptive textual and transcriptive verbal behavior are desirable initial repertoires for a BPDA. There has already been progress in these areas. Sejnowski and Rosenberg (1986) trained a neural network to read out loud. Other neural network researchers are making good progress on the problem of speech recognition and transcription. It will also eventually be desirable for PDAs to accept either spoken or typed or graphic input, and to produce printed or synthetic speech outputs (as well as graphics). It is conceivable that future generations of PDAs, with video and visual pattern recognition, may also be able to interpret gestures and facial expressions. With computer generated graphics, it is even possible to display a facial image that mimics human expressions. Stored video outputs or images constructed on-line may also supplement a language description of events.

Beyond these simple input and output media, however, a BPDA should both respond to and make use of idioms, slang, and autoclitics. The PDA which is dedicated to a single user will adapt very specifically to the verbal practices of that speaker/listener.

One of the most common and frustrating experiences for all computer users is to be unable to get the computer to do what they want, either because they do not know the precise command or because the desired behavior is simply not programmed. The computer is extremely rigid and precise. The verbal interface that will characterize a BPDA will provide a radically different experience, permitting the user to instruct, explain, teach, and ask to get the BPDA to do what is desired, like a good personal secretary. The BPDA will facilitate its own responsiveness to users by means of a skilled repertoire of asking for clarification, confirmation, examples, etc.

ELEMENTS OF VERBAL BEHAVIOR

For any system to acquire operant behavior, there must be primary reinforcers which strengthen its behavior. The BPDA should always act to maximize the

values of a human, i.e., what is reinforcing to the user should strengthen the behavior of the BPDA. While it is not currently feasible to link the computer directly into a human reinforcement process, the next-best solution may be for the human to "praise" behavior or outcomes that are desired. The human may also specify outcomes which should serve as reinforcers (e.g., profit produced by an electronic financial BPDA). These outcomes may be programmed ("genetically") as primary reinforcers, or their conditioned reinforcing value may be learned by the system based on their relation to user praise. We sometimes use the concept of *agent* in describing this solution to the reinforcer definition problem; i.e., a person or system who behaves in ways that promote the interest of its users or clients because they are the sole source of its primary reinforcers. This bridges the gap between those who want to think of a computer only as a tool and those who fear that computers with their own reinforcers are potentially a competitive threat to homo sapiens.

The simplest behavior of a BPDA as a speaker will be to simply tact that which can be computed, retrieved, or inferred. On the surface, this does not seem much different than what most computer systems do today. Currently, presentation of output is a mechanical act which does not require any verbal component. But in the BPDA, both the form and occurrence of output will be controlled by its effect on the system user as listener, and therefore the dynamics of this synthetic tact become more familiar to students of verbal behavior. Tacting by the BPDA will be strengthened either by explicit confirmation by the user that the information was useful or by evidence that effective action on the part of the user was successfully enabled. In order not to repetitively tact relationships which are of marginal value to their user, the BPDA must learn the value to the user of novel information or information relating most closely to the user's long-term or short-term objectives. Abstraction is an extremely important property of such

functional behavior, in marked contrast with most cognitive models.

Since the BPDA's reinforcers are very different from those of biological systems, its typical mands will be different. If the only primary reinforcer is user "praise," the only mand which specified a primary reinforcer would be asking the user for praise. However, in order to earn user praise, the BPDA will often execute sequences of actions, producing intermediate results which will have acquired conditioned reinforcing value. These conditioned reinforcers will presumably be the main object of mands. For example, the BPDA may mand the user for information, rules, suggestions, clarification of goals, etc. Fanciful as it might seem, a BPDA might ask for praise after accomplishing a goal for an unappreciative user, tactfully reminding the user of the functionality of specific contingent praise in shaping and maintaining its own most effective behavior.

Intraverbal responses, where a verbal stimulus to this net strengthens other verbal responses, will be very common in most BPDAs.

Human verbal behavior is currently distinguished from computer input and output on the basis of format. The computer response is typically very rigid and pre-specified, and human input is constrained by fixed requirements. However, it will be desirable for BPDAs to learn to interpret response forms by users that identify varying degrees of probability, frequency, or strength. The BPDA should also be able to emit these. These are, of course, the descriptive and quantifying autoclitics. Qualifying autoclitics will obviously be useful to the user in specifying relationships. Much research will need to be done on the value of the BPDA learning autoclitic frames, but it is likely that experience within the human verbal community will make these forms of BPDA response effective for many users. Word order is an important autoclitic function.

The verbal characteristics of the BPDA can only be imagined at this point, but much of the commercial value of these software agents will be determined by

their ability to communicate in a more human-like fashion with their users. But the BPDA might be of even more interest to the verbal behavior research community as a controlled laboratory for the dynamics of verbal relationships.

A VISION OF 7G COLLABORATION AND COMMERCIAL SUCCESS

Obviously BehavHeuristics, Inc. is interested in the long-term commercial potential of BPDAs. But as the rule-based Expert System did for cognitive science, a commercial success for BPDA also will bring favorable publicity and credibility to the entire behavioral community. The current business plan for BehavHeuristics includes the formation of an alliance between the academic behavioral community, private industry, and government funding agencies. The objectives will be twofold: to produce the 7G software to form the basis of the BPDA and to build up a library of training material. The training material will first need to enable the BPDA to acquire basic verbal behavior, and then to acquire specialized vocabulary and knowledge appropriate for specific problem domains. The 7G project will provide an ideal environment for laboratory research on fundamental principles of verbal behavior, and thus will serve the interests of basic research. Computer based training intended for humans could be debugged on an appropriately prepared BPDA. This training material will become part of the library of software that will be used to generate BPDAs. Once trained to a particular level or in a particular subject domain, BPDAs can easily be cloned.

We propose to proceed along the following lines. First, we hope this paper will stimulate the interest of leading researchers in verbal behavior. We encourage these researchers to contact us. A group of almost 100 behavior analysts interested in this general topic was organized by the second author in 1984; we propose to resurrect this network of people. Participants need to arrange communication mechanisms such as meetings at the Association for Behavior Analysis

and an electronic conference facilitated by "groupware." Given a mix of academic researchers wanting to publish and entrepreneurs wanting to earn a living, we need to establish mechanisms and rules to resolve the potential conflict between scientific openness and protection of intellectual property.

CONCLUSION

What we have proposed will be considered extremely ambitious by those in the field of artificial intelligence and neural networks who look at this task as an extension of Natural Language Processing and generative grammar, supposedly requiring much more complicated hardware than needed for nonverbal behaviors. In contrast, we believe that the development of synthetic verbal repertoires is an attainable engineering challenge, building upon the software capabilities that are already available and the cooperation of our colleagues who have specialized in the analysis of verbal behavior. This will provide an exciting research tool for the research community as well as practical applications for those who foresee a new generation of commercial personal digital assistants.

REFERENCES

- Bell, T. M., Hutchison, W. R., & Stephens, K. R. (1992). Using adaptive networks for resource allocation in changing environments. In P.J.G. Lisboa (Ed.), *Current Perspectives in Neural Networks*. London: Chapman and Hall.
- Cherpass, C. (1992). Natural language processing, pragmatics, and verbal behavior. *The Analysis of Verbal Behavior*, 10, 135-147.
- Catania, A. C., Shimoff, E., & Matthews, B. A. (1989). An experimental analysis of rule-governed behavior. In S. C. Hayes, (Ed.), *Rule-governed behavior: Cognition, contingencies, and instructional control* (pp. 119-150). New York: Plenum Press.
- Donahoe, J. W., & Palmer, D. C. (1989). The interpretation of complex human behavior: Some reactions to *Parallel Distributed Processing*, edited by J.L. McClelland, D.E. Rumelhart, and the PDP Research Group. *Journal of the Experimental Analysis of Behavior*, 51, 399-416.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine*. New York: The Free Press.
- Gupta, P., & MacWhinney, B. (1992). Integrating category acquisition with inflectional marking: A model of the German nominal system. In *Proceedings of the fourteenth annual conference of the cognitive science society* (pp. 253-264). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hinton, G. E. (1991). Mapping part-whole hierarchies into connectionist networks. In G.E. Hinton (Ed.), *Connectionist symbol processing*. Cambridge, MA: MIT Press.
- Hutchison, W. R. (1984, December). Cognitive vs. behavioral approaches to artificial intelligence. *Newsletter for the Behavioral Artificial Intelligence Network*, 1, 5-20.
- Hutchison, W. R. (1985a, March). Behavior analysis and other adaptive systems theories within the "functional camp" of artificial intelligence. *Newsletter for the Behavior Artificial Intelligence Network*, 2, 3-18.
- Hutchison, W. R. (1985b, May). Behavior analytic models for artificial intelligence. Paper presented at the annual meeting of the Association for Behavior Analysis, Columbus Ohio.
- Hutchison, W. R., & Stephens, K. R. (1986, May). Artificial intelligence: Demonstration of an operant approach. Paper presented at the meeting of the Association for Behavior Analysis, Milwaukee, WI.
- Hutchison, W. R., & Stephens, K. R. (1987a). Integration of distributed and symbolic knowledge representations. *Proceedings of the First International Conference on Neural Networks*, 2, 395-398. IEEE Press.
- Hutchison, W. R., & Stephens, K. R. (1987b). The Airline Marketing Tactician (AMT): A commercial application of adaptive networking. *Proceedings of the First International Conference on Neural Networks*, 4, 753-756. IEEE Press.
- Hutchison, W. R., & Stephens, K. R. (1988, May). Behavior Analytic Artificial Intelligence and "neural networks." Preconvention institute presented at the Association for Behavior Analysis, Philadelphia, PA.
- Hutchison, W. R., Stephens, K. R., & Bell, T. M. (1988, May). Behavioral implementations and applications of adaptive networks. Paper presented at the annual meeting of the Association for Behavior Analysis, Philadelphia, PA.
- Hutchison, W. R., & Stephens, K. R. (1989, February). The Airline Marketing Tactician: Allocation of resources in changing environments. Paper presented at the Second European Seminar on Neural Computing: Commercial Prospects, London, England.
- Hutchison, W. R., Stephens, K. R., & Bell, T. M. (1990). Adaptive networks for scheduling and inventory control: Resource allocation in changing environments. Paper presented at the Third European Seminar on Neural Computing: The Marketplace, London, England.
- Kehoe, E. J. (1989). Connectionist models of conditioning: A tutorial. *Journal of the Experimental Analysis of Behavior*, 52, 427-440.
- Klopf, A. H. (1982). *The hedonistic neuron: A theory of memory, learning, and intelligence*. Washington, D.C.: Hemisphere.
- Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tense of English verbs. In Rumelhart, D. E., McClelland, J. L., and the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Volume 2*. Cambridge, MA: MIT Press.
- Rumelhart, D. E., McClelland, J. L., & The PDP Research Group (Eds.), (1986). *Parallel distributed processing: Explorations in the microstructure of cognition, Volume 1*. Cambridge, MA: MIT Press.
- Schlinger, H. D. Intelligence: Real or Artificial? *The Analysis of Verbal Behavior*, 10, 125-133.

- St. John, M. (1992) Learning language in the service of a task. In *Proceedings of the fourteenth annual conference of the cognitive science society* (pp. 271-276). Hillsdale, NJ: Erlbaum.
- Sejnowski, T. J., & Rosenberg, C.R. (1986). NETtalk: a parallel network that learns to read aloud. Baltimore: The Johns Hopkins University Electrical Engineering and Computer Science Technical Report JHU/EECS-86/01. Reprinted in Anderson, J. A. & Rosenfeld, E. (Eds.) (1988). *Neurocomputing: Foundations of Research* (pp. 663-672). Cambridge, MA: The MIT Press.
- Skinner, B. F. (1957). *Verbal behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Skinner, B. F. (1969). *Contingencies of reinforcement: A theoretical analysis*. New York: Apple-Century-Crofts.
- Skinner, B. F. (1985). Cognitive science and behaviorism. *British Journal of Psychology*, 76, 291-301.
- Stephens, K. R., Hutchison, W. R., & Bell, T. M. (1988, May). Concepts, models and theories in adaptive networks. Paper presented at the annual meeting of the Association for Behavior Analysis, Philadelphia, PA.
- Stephens, K. R., & Hutchison, W. R. (1989). Behavior analysis and intelligent machines. Unpublished manuscript.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88, 135-171.
- Werbos, P. J. (1974). Beyond regression: New tools for prediction and analysis in the behavioral sciences. Unpublished Doctoral Dissertation, Harvard University.
- Werbos, P. J. (1977). Advanced forecasting methods for global crisis warning and models of intelligence. *General Systems Yearbook*, 22, 25-38.