Symbolic and Sub-symbolic Representations in Computational Models of Human Cognition

What Can be Learned from Biology?

Troy D. Kelley
ARMY RESEARCH LABORATORY

ABSTRACT. The debate over symbolic versus sub-symbolic representations of human cognition has been continuing for thirty years, with little indication of a resolution. The argument is this: Does the human cognitive system use symbols as a representation of knowledge, and does it process symbols and their respective constituents? Or does the human cognitive system use a distributed representation of knowledge, and is it somehow capable of processing this distributed representation of knowledge in a complex and meaningful way? This paper argues for an integrated symbolic and sub-symbolic approach to the representation of cognition. The lines of reasoning used as evidence to bolster this argument for an integrated approach are the cognitive architecture the Adaptive Character of Thought-Rational (ACT-R), and biology, where it is argued that symbolic and sub-symbolic representations of cognition are part of an intellectual continuum, with sub-symbolic representations at the low end and symbolic representations at the higher end.

KEY WORDS: evolutionary psychology, hybrid architectures, integrated cognitive architectures, knowledge representation

Symbolic and Sub-symbolic Representations in Cognition

The US Army Research Laboratory's (ARL) Human Research and Engineering Directorate (HRED) has begun an ambitious project to model human cognition on its high-performance computing (HPC) assets (a.k.a. supercomputers). The project, which is called Modeling and Integration of Neurological Dynamics with Symbolic Structures (MINDSS), has developed cooperative research agreements with major universities (Carnegie Mellon University, University of California, University of Massachusetts) and with other government organizations (National Aeronautics and Space Administration), with the goal of developing computational models of human

THEORY & PSYCHOLOGY Copyright © 2003 SAGE Publications. VOL. 13(6): 847–860
[0959-3543(200312)13:6;847–860;039293]

cognition which include the major areas of brain functionality (i.e. language, memory, perception) (Kelley, 2001). To achieve this goal, we have developed a theoretical perspective of human cognition that calls for the integration of the symbolic and sub-symbolic paradigms into a more cohesive theoretical whole. This paper presents the theoretical arguments for symbolic and sub-symbolic representations and then presents arguments for an integration of symbolic and sub-symbolic architectures.

Introduction

The question of how to computationally represent knowledge is a difficult one. The thesis that knowledge is represented as a system of symbols has been argued by classic cognitive psychology. Does the human cognitive system use symbols as a representation of knowledge, and process symbols and their respective constituents? Or does the human cognitive system use a distributed representation of knowledge, and is it somehow capable of processing this distributed representation of knowledge in a complex and meaningful way? The latter is the argument traditionally made by those in the artificial intelligence field (specifically, those in the connectionist movements).

This paper argues for an integrated symbolic and sub-symbolic approach to the computational representation of cognition. It presents three arguments as evidence to bolster the idea of an integrated approach to the development of cognitive architecture: First, the Adaptive Character of Thought-Rational (ACT-R) is a well-established cognitive architecture that includes aspects of both symbolic and sub-symbolic theory. Second, the biological perspective shows that symbolic and sub-symbolic representations of cognition are part of an intellectual continuum, with sub-symbolic representations at the low end and symbolic representations at the higher end. Third, the biological perspective will be discussed within the human anatomical system as well as across other species' anatomical composition.

Symbolic and Sub-symbolic Representation

The symbolic representation of human cognition can be easily understood if one uses the metaphor of a computer. Indeed, the development of the computer greatly influenced the symbolic approach to the study of cognition (Turing, 1950), an approach which has been called 'a species of computing, carried out in a particular type of biological mechanism' (Pylyshyn, 1989, p. 2). The function of a computer, in its most basic terms, can be thought of as an input$\Rightarrow$process$\Rightarrow$output system. The computer can take a series of symbols as input. These symbols are representations of some other concept or construct (which actually have meaning only to the human operator). The computer can then manipulate these symbols by using some pre-set instruction set. It can then output a result of the symbols based on the previous manipulation process. So, if a computer is given the number '4' and instructed to add the number '4' to the number '7', it will output the symbolic result '11'.

In the symbolic tradition, human cognition can be thought of in much the same way, as a symbolic manipulation process. If I am told, 'Mary loves Sam', I understand that 'Mary' is the symbolic representation for some person who loves Sam. I also understand that 'loves' is a symbolic representation of an emotional attachment two people can form for each other, and I understand that the symbol 'Sam' represents the person who is the receiver of Mary's affection. That is what I understand from the symbolic representation given to me. I can, however, infer other types of symbolic relationships, based on the one I was given. For example, I might infer that Sam is probably a male, and that Mary is probably a female. I might infer that Sam probably loves Mary as well. I might also infer that the relationship is probably serious, and so on. Note here that the problems involved with the inferring of information from symbolic relationships are extremely complex and difficult to implement computationally.

Sub-symbolic or connectionist systems are most generally associated with the metaphor of a neuron. Early implementations of sub-symbolic systems were called perceptrons (Rosenblatt, 1958, 1962). Like a small collection of neurons in the brain, a sub-symbolic system is composed of a small collection of perceptrons that operate in parallel to recognize a given input. This recognition process is accomplished by the adjustment of the weights which connect the perceptrons to each other. A collection of nodes can thus be enabled to recognize a given input and produce a specified output by adjusting the weights of the connections between the perceptrons. Therefore, a sub-symbolic system can be thought of as an autonomous learning system, and this is one of its great strengths.

Sub-symbolic architectures and symbolic architectures offer strengths and weaknesses to the study and the representation of human cognition. As previously mentioned, the sub-symbolic neural network can be viewed as an autonomous learning system, where a predetermined learning algorithm allows the network to relate input with appropriate output. Therefore, a sub-symbolic architecture can be 'trained' once care has been taken to develop the training set and develop the neural network architecture. Conversely, the symbolic architecture must have its internal representation of the world written by a knowledgeable programmer. Thus, the problem of representing the world can become a chore for a team of developers, as happened with the Cyc project (Lenat & Guha, 1989). Also, symbolic systems have a tendency to be removed from any real-world experience since their 'experience' is essentially provided by a programmer. This has been referred to as the symbol grounding problem (Harnad, 1989).

Perhaps the biggest problem facing either the symbolic or sub-symbolic architecture is the problem facing sub-symbolic architectures and their apparent difficulty in representing complex relationships. For example, a sub-symbolic architecture might be able to represent the constituent relationship 'Mary loves Sam' in terms of vector relationships, but more complex relationships such as 'Mary loves Sam and Bob hates Sally' tend to give sub-symbolic systems more trouble. A typical sub-symbolic model has difficulty making the distinction between 'Mary loves Sam and Bob hates Sally' and the slightly different 'Mary hates Sam and Bob loves Sally'. A sub-symbolic system can easily confuse these two symbolic relationships. This is a complex problem and is based on the constituent relationships among the symbols. As Fodor and Pylyshyn (1988) explain, 'It is important to see that this problem arises precisely because the [sub-symbolic] theory is trying to use sets of atomic representations to do a job that you really need complex representations for' (p. 322).

One can easily see the dichotomy between these two different types of approaches to the representation of cognition. Sub-symbolic systems operate in parallel—that is, many perceptrons can be used to recognize a given input—whereas symbolic systems operate in series—that is, symbolic representations are performed in a sequential manner. Sub-symbolic systems have a distributed representation of knowledge: that is, the symbolic representation of the relationship between '4 + 5' is not stored in any one location but is distributed across the various weights of the perceptrons. The symbolic system, by contrast, stores the symbol '4' in a given location in order to manipulate it with symbolic operations. Sub-symbolic systems learn to recognize input and respond in accordance with a learning rule. However, symbolic systems are concerned not with the recognition of a stimulus, but, instead, with the manipulation of recognized symbols following the recognition. While there is a dichotomy between these two approaches to cognition, one can also view the two approaches as two ends of a single continuum, especially when one thinks of the continuum in these terms: sub-symbolic systems recognize input and pass that input along to more symbolic systems.

Integrated Architectures

The cognitive architecture ACT-R (Anderson & Lebiere, 1998) represents an interesting integration of symbolic and sub-symbolic mechanisms. As background, the ACT-R architecture is a symbolic, production system architecture, capable of low-level representations of memory structures. Production system architectures are those where the main type of processing occurs within an If–Then format. It includes a declarative memory component and a procedural memory component. Declarative memories are those which can be characterized as long-term factual memories (i.e. phone numbers, dates, locations). Procedural memories can be characterized as low-level compilations of memories, primarily of basic skills, which are sometimes difficult to describe to other people (how to ride a bike, how to shoot a basketball). ACT-R is implemented in the common LISP programming language as a collection of LISP functions and sub-routines, which can be accessed by the cognitive modeler. More importantly, ACT-R has an underlying layer of sub-symbolic processes that affect the higher layers of symbolic processes. This does not mean that ACT-R is constructed on top of a neural network or some other type of connectionist framework. Instead, the symbolic components of ACT-R are linked to the sub-symbolic framework through a series of equations that 'determine many of the continuously varying, qualitative properties of the symbolic cognitive elements' (Anderson & Lebiere, 1998, p. 13). In practical terms, this means that if I wanted to use ACT-R to represent the cognition involved with remembering the symbol '4', I could attach the symbol '4' to a decay function; this function would simulate, in a continuously varying manner, the decay and possible forgetting of the symbol '4'. Furthermore, if I needed to remember a series of numbers, each of the memories for these numbers would have a continuously varying algorithm associated with them; and these algorithms would vary in parallel with each other. Thus, in ACT-R, sub-symbolic algorithms are able to vary continuously and in parallel while symbolic operations continue in serial fashion. This is just one example of an integrated cognitive architecture.

While ACT-R can function as an integrated symbolic–sub-symbolic architecture, it still is prone to the deficiencies of each architecture. For example, much of the knowledge that is acquired within ACT-R is written by an experienced programmer and not developed through recognition of outside stimuli. In other words, ACT-R does not use sub-symbolic mechanisms to recognize and identify stimuli; instead, knowledge structures which relate to the outside world are developed by a programmer. ACT-R can create new declarative memory structures as it processes information, but these are generally not perceptual aspects of the environment (this distinction is being made of ACT-R 4.0). With the development of ACT-R 5.0, newer perceptual-motor components have allowed for some integration of perceptual and motor elements with declarative memory structures. However, these newer systems are not of the type capable of developing a learned recognition of external stimuli, as is typical of sub-symbolic systems. Nevertheless, ACT-R does show some of the advantages of an integrated hybrid architecture, and has set a trend within applied psychology toward more integrated cognitive architectures.

Hybrid architectures are becoming more and more popular as a way of addressing the limitations of purely symbolic or sub-symbolic approaches. Sun and Bookman (1995) and Sun and Alexandre (1997) review the state-of-the-art in hybrid architectures and offer many examples of successful sub-symbolic and symbolic integration approaches. While many of the papers in the latter volume concentrate on the more philosophical architectural issues of integration (i.e. how much integration, what type of integration, direction of information flow), there are also examples of successfully integrated architectures. One of the more successful implementations was the CLARION architecture (Sun & Peterson, 1997), which was able to improve over traditional sub-symbolic learning techniques (Q-learning) and successfully generalize learned rules to other domains.

The Human Nervous System

This paper puts forth a biological argument to support the view that cognition should be represented as a symbolic–sub-symbolic integrated system, one that uses sub-symbolic input at its lowest end and manipulates symbolic representations at its highest end in order to produce appropriate output. To examine this argument, the biology of cognition will be approached from two different angles: first, the biology and capability of the human cognitive system will be specifically examined; then the biology and capability of cognitive systems across a continuum of species will be examined. From these two points of view, the hope is to build the evolutionary argument for sub-symbolic and symbolic integration.

The internal elements of the human cognitive system can be viewed as two ends of a cognitive continuum. At the highest end, the symbolic processing end, the human system is composed of the prefrontal cortex areas of the brain. As a structural component within the human brain, the prefrontal cortex represents the zenith of human cognitive capabilities. Early neuro-psychologists labeled the frontal lobes as the 'seat of wisdom' (Beaumont, Kenealy, & Rogers, 1996, p 348). The prefrontal cortex and the larger, more encompassing frontal lobes play an important role in many of the most complex of human activities. Behaviors such as self-awareness, planning and supervisory control emanate from the frontal lobes. It is in these areas of the brain that human reasoning of the most complex and symbolic form originates. Furthermore, these systems have evolved on top of lower, older, more simple systems. Within the frontal lobes, the neocortex is primarily a mammalian structure that is built upon the older and more primitive structures of the human brain (Marin-Padilla, 1988).

Language represents a symbolic system. The proponents of the symbolic representations of human cognition often characterize human reasoning as the 'Language of Thought' (Fodor, 1975; Pylyshyn, 1984). It is clear that the human cognitive system is capable of symbolic manipulation and much of this activity occurs in the frontal lobe of the brain or the high end of the human cognitive spectrum. The frontal lobes, more specifically, the prefrontal cortex, play a particularly large role in the processing of language. Damage to the prefrontal cortex can lead to various language impairments including dysarthria, aphemia and Broca's Aphasia (Beaumont et al., 1996). While some of human language processing also occurs in Wernicke's area, which is located in the temporal lobe, it is important to note that Wernicke's area is also part of the neocortex, which is the outermost layer of the brain. Researchers have found that the bulk of language processing is cortically based (Metter et al., 1988).

At the lowest end of the human cognitive spectrum, the simplest cognitive mechanism is the reflex. Simple reflexes are the most basic neural architecture, since only one synapse is involved in order to elicit an appropriate response. For example, the myotatic stretch reflex (or the tendon reflex), which has the effect of stabilizing a muscle in response to an external flex, happens 'quite quickly and unconsciously because only one synapse is involved' (Matsumoto, Walker, Walker, & Hughes, 1990, p. 94). The function of a reflex (at least the most simple of reflexes) is to remove any input from the higher cognitive functions of the brain. In other words, the sub-symbolic system of reflexes is devoid of any symbolic processing. This has two advantages: (1) to accelerate the response to the stimulus; and (2) to remove any unneeded cognition from the decision process. The analogy to the sub-symbolic system is that the reflex, especially a myotatic stretch reflex, could be thought of as a single feed-forward neural network. The reflex has been pre-trained through millions of years of evolution to recognize a specific input, and the appropriate output is then emitted; no interference from higher cognitive elements is needed. To make the analogy to a sub-symbolic system, the weights of the neural net have been set by evolution.

Biological Architectures

Just as within the human anatomical system there is an intellectual continuum from the simple reflex to the frontal lobes, so too in the animal kingdom there is an intellectual continuum which has evolved over time from the lowest organisms to the highest organisms. At the bottom of the evolutionary tree are simple one-celled organisms which have extremely rudimentary nervous systems. However, these organisms are capable of displaying elementary behaviors such as avoidance of heat, cold or noxious chemicals. All organisms at the bottom of the evolutionary tree represent a class of organisms displaying the simplest of nervous systems, which could then be replicated by the simplest of sub-symbolic models. The basic two-layered sub-symbolic models developed during the 1940s by McCullough and Pitts (1943) were capable of reacting in ways similar to the behaviors of single-celled organisms.

Continuing to move up the evolutionary tree, simple organisms begin to develop simple nervous systems also called ‘nerve nets’. The neurons in an organism called the hydra display some of the same anatomical characteristics of human neurons, including communications via a gap between neurons called a synapse. As complexity increases, neurological organization begins to include some resemblance to a spinal column and a primitive brain. For example, some organisms have long nerve cords which are connected to cerebral ganglia located in the head. These organisms also have sensory capabilities called auricles that are concentrated in the head region. These auricles contain chemical receptors that are used to gather food; in addition, they have structures similar to eyes, called ocelli, which help them to avoid light. The point is that a ‘nervous system complexity continuum’ starts to emerge as one progresses up the evolutionary tree from invertebrates to vertebrates, from organisms that can respond to the simplest of stimuli with some kind of distributed nervous system to organisms that have some type of primitive sensory system that allows them to locate food. The progression continues from organisms that have small brain-like structures to organisms that are capable of associative learning.

The learning capabilities of insects are frequently described as ‘associative learning’ (Simpson & White, 1990). Associative learning is the simple association of a stimulus with a response. Interestingly, early sub-symbolic models were actually criticized as being capable of exhibiting *only* associative learning. The main point of the criticism was that simple association was not capable of learning more complex relationships (Fodor & Pylyshyn, 1988). This critique is similar to the assertion that atomic representations are not detailed enough to represent complex relationships. More complex ways of representing knowledge and processing knowledge are needed by higher-level organisms in order to process symbolic relationships. And as the progression to the higher levels of nervous system complexity continues, one begins to see the ability of the higher-level organisms to process not just simple associations, but symbolic relationships as well. This symbolic representational understanding is evidenced by the understanding of language by highest-level organisms, namely apes and humans.

There is strong evidence that apes are capable of understanding symbolic relationships similar to those used in human speech. As early as the 1950s, researchers found that an ape named Viki was capable of sorting photographs into conceptual categories (Hayes & Hayes, 1953). A few decades later, a chimpanzee named Washoe was taught ASL (American sign language) and was able to produce 85 different signs by the age of 36 months (Gardner & Gardner, 1971). Besides Washoe's use of ASL, other types of symbolic systems have been used to teach chimpanzees to 'talk'. Rumbaugh et al. (1973) were able to teach chimpanzees to use non-verbal symbolic systems, which were essentially symbols that could be selected via buttons on a screen, which could then be combined to produce sentence-like structures. Rumbaugh et al.'s work was later replicated by researchers in Japan using a similar symbolic representational system (Matsuzawa, 1989; paraphrased from Savage-Rumbaugh et al., in press).

The use of symbols by apes has been criticized by some researchers (Terrace, Straub, Bever, & Seidenberg, 1977: Thompson & Church, 1980) as being the learning of simple associative chains, similar to what pigeons could learn. Others (Sebeok & Umiker-Sebeok, 1980) have considered the symbol manipulation of apes to be an elaborate form of cueing. All this criticism has been addressed and shown to be unsubstantiated (Rumbaugh, 1981; Rumbaugh & Savage-Rumbaugh, 1980; Savage-Rumbaugh, 1986, 1987; Savage-Rumbaugh, Romski, Sevcik, & Pate, 1983; Savage-Rumbaugh & Rumbaugh, 1982).

Furthermore, an argument could be made that the apes are not actually understanding symbols at all, but instead simply translating symbols without an understanding of their true meanings. This is the famous 'Chinese Room Argument' (Searle, 1979). In the Chinese Room Argument, Searle postulates that a person using a look-up-table could translate Chinese into English with little or no understanding of Chinese. However, the Chinese Room Argument has been criticized as not addressing the symbol-grounding problem (Harnad, 1989). The symbol-grounding problem acknowledges that many symbols are grounded by perception and experience. Apes being tested on language comprehension are not exposed to the symbol-grounding problems of computer simulations of the mind because they have experience in the world which computers do not. Tests of ape symbolic comprehension could easily be addressed by tests not of only symbolic comprehension, but also of the actual meaning of symbols. These tests have been done, and the most recent research by Savage-Rumbaugh and her colleagues concludes that apes can comprehend symbols (Savage-Rumbaugh et al., in press).

There is biological evidence for an intellectual continuum within the human cognitive system as well as within the biological system of species as a whole. At one end of the human spectrum, reflexes approach the simplest of feed-forward sub-symbolic systems. Likewise, at the low end of the evolutionary scale, simple organisms display avoidance behaviors that could be easily simulated by sub-symbolic models. At the higher end of the biological continuum, apes and humans appear to be able to manipulate symbolic systems and communicate via manipulation of symbols. It is also possible that other higher-order organisms, such as dolphins, are capable of symbolic manipulations as well (Herman, Matus, Herman, & Pack, 2001).

Conclusions

The idea that symbolic and sub-symbolic architectures should be integrated is not a new one. Researchers have been working for the past few years on integrated architectures and have produced a wealth of research (Barnden & Holyoak, 1994; Gallant, 1993; Goonatilake & Khebbal, 1995; Medsker, 1994, 1995; Miikkulainen, 1993; Reilly & Sharkey, 1992; Smolensky, 1995, Sun & Alexandre, 1995, 1997; Wermeter, Riloff, & Scheler, 1997). However, the reasoning for the integration has never been made on biological grounds. Psychological plausibility has been a central theme of both symbolic researchers and sub-symbolic researchers, and yet it has never been used as justification for an integration of architectures. Psychological plausibility has long been one of the hallmarks of sub-symbolic systems with their albeit simple representations of interconnected neurons that are capable of learning. Recently, symbolic architectures have been striving to reach some level of congruence with psychological plausibility as well. For example, using Functional Magnetic Resonance Imaging (*f*MRI), the hybrid symbolic architecture ACT-R has shown some correlations with the cognitive mechanisms revealed by FMRI technology (Sohn, Ursu, Anderson, Stenger, & Carter, 2000). On the sub-symbolic side of pure connectionist models, newer sub-symbolic implementations have emphasized their correlations with psychological plausibility as well (Shastri & Ajjanagadde, 1993).

It could be argued that symbolic architectures do not need to be placed at the top of the intellectual continuum, instead that sub-symbolic systems merely need to catch up to symbolic system's level of representational complexity. While it has been shown that the traditional connectionist typology is inadequate for representing complex symbolic relationships (Fodor & Pylyshyn, 1988), newer, more complex sub-symbolic topologies may allow for a strictly sub-symbolic representation of the entire mind. For example, Smolensky (1995) has used the tensor product of vector relationships in order to represent recursion and thus overcome some of the problems associated with sub-symbolic representation of constituent relationships. Also, Shastri and Ajjangadde (1993) has used dynamic sub-symbolic systems where the synchronous firing of related neural networks is used to represent conceptual relationships. Shastri and Ajjangadde's pulsed neural nets also address other limitations of sub-symbolic systems, namely that they tend to be static representations of some previously presented pattern of inputs, while the brain's behavior is better characterized as a dynamic system of electro-chemical pulses. While it has been noted that the static representations of traditional neural networks could be viewed as very similar to a reflex response; clearly more complex dynamic designs are evident in the frontal lobes of the brain. It could be that further research into different types of sub-symbolic topologies will led the representational complexity needed for a fully sub-symbolic representation of the mind.

While increased complexity of sub-symbolic architectures may be one answer to adopting a complete sub-symbolic architecture for the representation of cognition, the traditional symbolic architecture still holds numerous advantages for the representation of *complex* cognition. Until one architecture can clearly display an overwhelming advantage, researchers with the HRED MINDSS program will continue to investigate avenues for the integration of symbolic and sub-symbolic architectures. The argument put forth in this paper is that biology has evolved compelling examples of integrated architectures; both within human cognition and across species, there is evidence of an intellectual continuum of symbolic and sub-symbolic integration.

References

Anderson, J.R., & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Lawrence Erlbaum Associates,

Barnden, J.A., & Holyoak, K.J. (Eds.). (1994). *Advances in connectionist and neural computation theory: Vol. 3. Analogy, metaphor and reminding.* Norwood, NJ: Ablex.

Beaumont, J.G., Kenealy, P.M., & Rogers, M.J.C. (Eds.). (1996). *The Blackwell dictionary of neuropsychology*. Malden, MA: Blackwell.

Fodor, J.A. (1975). *The language of thought*. New York: Crowell.

Fodor, J.A., & Pylyshyn, Z.W. (1988). Connectionism and the cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.

Gallant, S.I. (1993). *Neural network learning and expert systems*. Cambridge, MA: MIT Press.

Gardner, B.T., & Gardner, R.A. (1971). Two way communication with an infant chimpanzee. In A.M. Schrier & F. Stollnitz (Eds.), *Behavior of nonhuman primates* (Vol. 4, pp. 117–135). New York: Academic Press.

Goonatilake, S., & Khebbal, S. (1995). *Intelligent hybrid systems.* Chichester: Wiley.

Harnad, S. (1989). Minds, machines and Searle. *Journal of Theoretical and Experimental Artificial Intelligence*, *1*, 5–25.

Hayes, K.J., & Hayes, C. (1953). Picture perception in a home-raised chimpanzee. *Journal of Comparative and Physiological Psychology*, *46*, 470–474.

Herman, L.M., Matus, E.Y.K., Herman, M.I., & Pack, A.A. (2001). The bottlenosed dolphin's (*Tursiops truncatus*) understanding of gestures as symbolic representations of its body parts. *Animal Learning and Behavior*, *29*, 250–264.

Kelley, T.D. (2001). An overview of modeling and integration of neurological dynamics with symbolic structures (MINDSS). In W.F. Waite (Ed.), *Proceedings of 2000 Summer Computer Simulation Conference.* (pp. 730–735). Society for Computer Simulation, San Diego, CA.

Lenat, D.B., & Guha, R.V. (1989). *Building large knowledge-based systems: Representation and inference in the Cyc project.* New York: Addison Wesley Longman.

Marin-Padilla, M. (1988). Early ontogensis of the human cerebral cortex. In A. Peters & E.G. Jones (Eds.), *Cerebral cortex* (Vol. 7, pp. 1–34). New York: Plenum.

Matsumoto, R., Walker, B., Walker, J.M., & Hughes, H.C. (1990). Fundamentals of neuroscience. In J.T. Cacioppo & L.G. Tassinary (Eds.), *Principles of psychophysiology* (pp. 58–112). New York: Cambridge University Press.

Matsuzawa, T. (1989). Spontaneous pattern construction in a chimpanzee. In P. Heltne & L.A. Marquardt (Eds.), *Understanding chimpanzees* (pp. 252–265). Cambridge, MA: Harvard University Press.

McCullogh, W.S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, *5*, 115–133.

Medsker, L.R. (1994). *Hybrid neural networks and expert systems.* Boston, MA: Kluwer Academic.

Medsker, L.R. (1995). *Hybrid intelligent systems*. Boston, MA: Kluwer Academic.

Metter, E.J., Riege, W.H., Hanson, W.R., Jackson, C.A., Kempler, D., & VanLancker, D. (1988). Subcortical structures in aphasia: An analysis based on (F-18)—florodoxyglucose positron emission tomography, and computed tomography. *Archives of Neurology*, *45*, 1229–1234.

Miikkulainen, R. (1993). *Subsymbolic natural language processing.* Cambridge, MA: MIT Press,

Pylyshyn, Z. (1984). *Computation and cognition: Toward a foundation for cognitive science.* Cambridge, MA: MIT Press/Bradford.

Pylyshyn, Z. (1989). Computing in cognitive science. In M. Posner (Ed.), *Foundations of cognitive science* (pp. 51–91). Cambridge, MA: MIT Press.

Reilly, R.G., & Sharkey, N.E. (1992). *Connectionist approaches to natural language processing.* Hillsdale, NJ: Erlbaum.

Rosenblatt, F. (1958). The perceptron: A probabilistic method for information storage and organization in the brain. *Psychological Review*, *65*, 386–408.

Rosenblatt, F. (1962). *Principles of neurodynamics.* New York: Spartan.

Rumbaugh, D.M., von Glasersfeld, E., Warner, H., Pisani, P.P., Gill, T.V., Brown, J.V., & Bell, C.L. (1973). A computer-controlled language training system for investigating the language skills of young apes. *Behavioral Research Methods and Instrumentation*, 5, 385–392.

Rumbaugh, D.M. (1981). Who feeds Clever Hans? In T.A. Sebeok & R. Rosenthal (Eds.), The Clever Hans phenomenon: Communication with horses, whales, apes and people [Special Issue]. *Annals of the New York Academy of Sciences*, *364*, 26–34.

Rumbaugh, D.M., & Savage-Rumbaugh, E.S. (1980, Winter). A response to Herbert Terrace's article 'Linguistic apes: What are they saying?' *New York University Education Quarterly*, p. 33.

Savage-Rumbaugh, E.S. (1986). *Ape language: From conditioned response to symbol.* New York: Columbia University Press.

Savage-Rumbaugh, E.S. (1987). Communication, symbolic communication, and language: Reply to Seidenberg and Petitto. *Journal of Experimental Psychology: General*, *116*, 288–292.

Savage-Rumbaugh, E.S., Murphy, J., Sevcik, R.A., Rumbaugh, D.M., Brakke, K.E. & Williams, S. (in press). Language comprehension in ape and child. *Monographs of Society for Child Development.*

Savage-Rumbaugh, E.S., Romski, M.A., Sevcik, R.A., & Pate, J.L. (1983). Assessing symbol usage versus symbol competency. *Journal of Experimental Psychology, General*, *112*, 508–512.

Savage-Rumbaugh, E.S., & Rumbaugh, D.M. (1982). Ape language is alive and well. *Anthropos, 77*, 568–573.

Searle, J.R. (1979). What is an intentional state? *Mind, 88*, 72–94.

Sebeok, T.A., & Umiker-Sebeok, J. (1980). *Speaking of apes: A critical analysis of two-way communication with man*. New York: Plenum.

Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, *16*, 417–494.

Simpson, S.J., & White, P.R. (1990). Associative learning and locusts feeding: Evidence for a 'learned hunger' for protein. *Animal Behavior*, *40*, 506–513.

Smolensky, P. (1995). Reply: Constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. In C. Macdonald & G. Macdonald (Eds.), *Connectionism: Debates on psychological explanation* (Vol. 2, pp. 223–290). Oxford/Cambridge, MA: Blackwell.

Sohn, M., Ursu, S., Anderson. J.R., Stenger, A.V., & Carter, C.S. (2000). The role of prefrontal cortex and posterior parietal cortex in task switching. *Inaugural Articles of the National Academy of Sciences, PNAS*, *97*(24), 13448–13453.

Sun, R., & Alexandre, F. (1995). *Proceedings of the workshop on connectionist-symbolic integration: From unified to hybrid approaches*. Montreal: McGraw-Hill.

Sun, R., & Alexandre, F. (Eds.). (1997). *Connectionist symbolic integration: From unified to hybrid approaches.* Hillsdale, NJ: Erlbaum.

Sun, R., & Bookman, L.A. (1995). *Computational architectures integrating neural and symbolic processes: A perspective on the state of the art*. Norwell, MA: Kluwer Academic Publishers.

Sun, R., & Peterson, T. (1997). A hybrid agent architecture for reactive sequential decision making. In R. Sun & F. Alexandre (Eds.), *Connectionist symbolic integration: From unified to hybrid approaches* (pp. 113–138). Hillsdale, NJ: Erlbaum.

Terrace, H.S., Straub, R.O., Bever. T.G., & Seidenberg, M.S. (1977). Representation of a sequence by a pigeon. *Bulletin of the Psychonomic Society*, *10*, 269.

Thompson, C.R., & Church, R.M. (1980). An explanation of the language of chimpanzee. *Science, 208*, 313–314.

Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, *59*, 433–460.

Wermeter, S., Riloff, E., & Scheler, G. (1997). *Connectionist, statistical and symbolic approaches to learning for natural language processing.* Berlin: Springer.

TROY D. KELLEY is an Engineering Psychologist at the Human Research and Engineering Directorate (HRED) of the Army Research Laboratory (ARL) at Aberdeen Proving Ground, MD. His recent work has concentrated on the understanding of human cognition, especially in areas of problem solving, strategy selection and decision-making. For the Army, he worked extensively with the modeling and analysis tool IMPRINT

(Improved Performance Research Integration Tool), which has been transitioned to over 120 sites throughout the Department of Defense R&D community. In 1998 he obtained a second Master's Degree from George Mason University in Human Factors/Applied Cognition, where he developed cognitive models using the ACT-R (Atomic Components of Thought-Rational) architecture. Currently he is heading a program called MINDS (Modeling and Integration of Neurological Dynamics with Symbolic Structures), which has the goal of modeling human cognitive processes on the Army Research Laboratory's HPC (high-performance computing) assets in cooperation with major US universities. ADDRESS: Army Research Laboratory, AMSRL-HR-SE, Aberdeen Proving Ground, MD, 21005-5425, USA. [email: tkelley@arl.army.mil]