Intelligence: Real or Artificial?

Henry D. Schlinger, Jr.
Western New England College

Throughout the history of the artificial intelligence movement, researchers have strived to create computers that could simulate general human intelligence. This paper argues that workers in artificial intelligence have failed to achieve this goal because they adopted the wrong model of human behavior and intelligence, namely a cognitive essentialist model with origins in the traditional philosophies of natural intelligence. An analysis of the word "intelligence" suggests that it originally referred to behavior-environment relations and not to inferred internal structures and processes. It is concluded that if workers in artificial intelligence are to succeed in their general goal, then they must design machines that are adaptive, that is, that can learn. Thus, artificial intelligence researchers must discard their essentialist model of natural intelligence and adopt a selectionist model instead. Such a strategic change should lead them to the science of behavior analysis.

In the early days of the Artificial Intelligence (AI) movement its proponents and participants set forth the lofty goal of simulating general human intelligence using computers. The consensus now, even among many AI researchers (e.g., Todd, 1992), is that they have failed to achieve this goal although many far less reaching, but still impressive, results have been attained. Consequently, AI researchers have been forced to revise their estimates of what they could accomplish. Perhaps the primary reason for the failure to simulate what we speak of as general intelligence is that the theoretical model that forms the basis of most AI research, which is derived from traditional philosophies of natural intelligence, is flawed. Obviously, if the assumptions upon which AI is based are wrong, then computer simulations following from those assumptions will fail. It is not that the lofty goal of simulating general human intelligence — whatever that is — is impossible, only that it will be impos-

sible so long as the theoretical model of AI is flawed. The position of this paper is that only a scientific functional analysis of human behavior can yield a successful working model from which computer modeling could then proceed.

A LITTLE HISTORY

The formal beginning of what came to be known as artificial intelligence was in 1956 when a group of mostly young mathematicians gathered at a summer institute at Dartmouth College for the purpose of discussing the possibilities of producing computer programs that could behave or think intelligently. Specifically, their discussions were fueled by the assumption that, as McCorduck stated, "every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it" (1979, p. 93, as cited in Gardner, 1985). This was indeed a lofty goal. The term "artificial intelligence" is widely believed to have been coined by John McCarthy, who was then an assistant professor of mathematics at Dartmouth and a major organizer of the conference. According to Howard Gardner, in his book *The Mind's New Science* (1985), most people in AI agree on

I am grateful to Dennis Kolodziejski for his helpful comments on an earlier version of this paper, which was presented in November, 1992 as part of the Faculty Forum series at Western New England College titled "Artificial and Natural Intelligence: Differences and Implications." Correspondence concerning this article should be addressed to the author, Department of Psychology, Western New England College, 1215 Wilbraham Road, Springfield, MA 01119.

two things. The first is that AI "seeks to produce, on a computer, a pattern of output that would be considered intelligent if displayed by a human being" (1985, p. 11), and the second is that AI is viewed as a way of testing a particular theory of how cognitive processes might work. That theory is the popular information-processing model of cognition. Where AI researchers disagree, according to Gardner, is how literally to interpret the thinking metaphor. For example, some take what John Searle calls the "weak view" of AI, wherein computer programs are simply a means for testing theories of how humans might carry out cognitive operations. The weak view of AI is synonymous with modern cognitive psychology. John Searle (1984) summarizes the cognitive position as follows:

Thinking is processing information, but information processing is just symbol manipulation. Computers do symbol manipulation. So the best way to study thinking . . . is to study computational symbol-manipulating programs, whether they are in computers or in brains . . . [T]he task of cognitive science is to characterize the brain, not at the level of nerve cells, nor at the level of conscious mental states, but rather at the level of its functioning as an information processing system. (p. 43)

Others, however, hold the view of a "strong AI," in which "the appropriately programmed computer really is a mind in the sense that computers given the right programs can be literally said to understand and have other cognitive states. In strong AI, because the programmed computer has cognitive states, the programs are not merely tools that enable us to test psychological explanations; rather the programs are themselves the explanations" (Searle, 1980, p. 417, cited in Gardner, 1985). Simply speaking, in strong AI the brain really is a digital computer and the mind really is a computer program (Searle, 1984).

Four Phases of AI Research

Hubert Dreyfus (1979) has described research in AI over a 20-year period from 1957 to 1977 according to four phases. Phase I, from 1957 until 1962, was characterized by programs in so-called cognitive simulation such as language translation, problem solving, and pattern recognition.

Examples of fairly successful programs in cognitive simulation included the Logic Theorist devised by Newell, Shaw, and Simon, which successfully proved 38 of 52 theorems from *Principia Mathematica*, and the General Problem Solver (GPS), which was capable of solving theorems, playing chess, and solving such problems as the missionary-cannibal conundrum, the tower of Hanoi, and cryptarithmic. Despite the fact that the GPS could perform these tasks "intelligently," according to Dreyfus, contrary to its name, it was not very general. Phase II, from 1962 to 1967, was characterized by semantic information processing. The objective during this phase "was to create a computer simulation of English language understanding with 'understanding' narrowly defined to mean that with a limited subset of English as input, the computer would respond with a limited subset of appropriate English" (Crossman, 1985). Again, attempts to generalize these programs met with failure (Dreyfus, 1979). Phase III, from 1967 to 1972 was characterized by microworlds, or computer programs with greatly restricted domains. Examples included, among others, DENDRAL, which successfully competed with human chemists by inferring chemical structures on the basis of a vast amount of data from mass spectrographs. Perhaps the most well-known example is Terry Winograd's program SHRDLU, which allowed a person to "engage in a dialogue with the computer" by using simple English commands which caused a simulated robot arm to manipulate a set of variously shaped blocks. Phase IV, from 1972 to 1977, which Dreyfus refers to as Knowledge Representation or Cognitive Science, was characterized by so-called "expert systems." A well-known example is MYCIN, a program for diagnosing blood and meningitis infections, and which did a respectable job of simulating the expertise of a medical practitioner.

THE FAILURE OF AI

Behavior analysts should have no problem agreeing with Hubert Dreyfus that both the weak and the strong view of AI

have failed. Specifically, AI failed both in its more modest attempt to support an information processing theory of human intelligence and in its more presumptuous attempt to create computer programs that behave in ways that lead us to say that they actually "understand." The first, and most pervasive criticism of AI is that the programs have not demonstrated the generality that characterizes the behavior of most "intelligent" organisms. In fact, critics like Dreyfus claim that these programs are merely engineering tricks which implies that the programs will continue to be successful only so long as they are limited to their very restricted domains. According to Dreyfus (1979),

The nongeneralizable character of the programs ... makes them engineering feats, not steps toward generally intelligent systems, and they are, therefore not at all promising as contributions to psychology. (p. 18)

Furthermore, Dreyfus (1979) argues that AI has failed in its goals because its workers have based their efforts on four assumptions that he believes are fundamentally incorrect. The biological assumption is that the brain and the digital computer are functionally similar in that they both process information digitally and serially. The psychological assumption is that even though the analogy of the brain as a computer might be weak, humans do process discrete bits of information like computers do. The epistemological assumption is that even if humans do not process information like computers do it may still be possible to formalize human behavior according to a different set of rules that can be reproduced by a computer. And finally, the ontological assumption is that if the world could be analyzed into enough context free atomistic data or facts, then intelligent behavior could be simulated. For Dreyfus, then, the fundamental metaphysical assumption behind most AI work is that both human behavior and the context in which it occurs can be analyzed according to a set of formal rules that could be written into a computer program to simulate intelligent behavior. It is thus not surprising to discover that "the theoretical foundations of modern cognitive science

are rooted in logic and computation theory" (Slack, 1984).

So far the metaphysical assumption of AI remains just that, a metaphysical assumption. Moreover, Dreyfus argues that AI has failed and will continue to fail so long as workers in the field ignore some important differences between computers and humans, namely that humans have bodies that sense the environment and act upon the environment, and most importantly, that a body has needs and that much if not all human behavior, like the behavior of all animals, involves the drive to survive. These last few criticisms can be summarized as a failure on the part of AI researchers to understand the nature of Darwinian selection not only at the level of genes but, more importantly, at the level of the behavior of individual organisms in their environment. The metaphysical assumption of AI ignores the functional aspects of an organism's adaptive behavior and, instead, stresses a logical formal structural approach that has no scientific support (Dreyfus, 1979). The point to be made is that the essence of human intelligence is not an ability to manipulate symbols or representations as cognitive scientists would have us believe, but rather the modifiability of behavior (Futuyma, 1979).

A FUNCTIONAL PERSPECTIVE: ESSENTIALISM VERSUS SELECTIONISM

The major implication of the above discussion of the failure of AI is that its workers have largely neglected the functional aspects of human intelligence. The remainder of this essay will focus on a functional analysis not only of behavior that we call intelligent, but also of the language we use to describe human intelligence. Such an analysis naturally raises several interesting philosophical questions. For example, What is intelligence? What is thinking? What is understanding? What is the best way to talk about human intelligence? Can human intelligence be simulated on a computer? However, to even ask, "What is intelligence?" is to fall into the essentialism trap. Phrasing questions in such ways

assumes that there is a property of human nature — called intelligence — with essential immutable qualities. For cognitive scientists who are interested in AI the task is to attempt to represent this quality in a machine. A functional, selectionist, approach, however, follows a different path. Rather than imposing *a priori* an essentialistic definition of intelligence and then looking for it as cognitive psychologists have a tendency to do, or representing it, as workers in AI attempt to do, researchers must *discover* the definition (see Palmer and Donahoe, 1992). This is to be done by conducting a functional analysis of the behavior they are interested in and looking for order therein.

Perhaps the most important question for the purpose of this essay and, hence, for the AI movement is, "What do we mean when we speak of intelligence?" Following from the title of this essay, we must ask whether intelligence, as it is usually conceptualized, is real or artificial? In other words, does the term "intelligence" refer to a qualitatively unique structure or process that can be defined *a priori*? The word intelligence comes from the Latin *intellegere*, meaning to perceive or understand, from the roots *inter* meaning between or among, and *legere*, meaning to gather, pick or choose. (Notice that these roots do not refer to inferred mental or cognitive processes occurring at an unobservable level of analysis, but rather to behavior — gathering, picking, and choosing — occurring in the environment.) Thus, intelligence or, more aptly, intelligent behavior is what we observe when we say that an individual perceives relationships or differences between or among situations. The fundamental question then, especially for AI researchers who are interested in generating intelligent systems, is how does the behavior that we call intelligent come about? In other words, how does it originate? How is it that individuals come to perceive the relationships or differences between or among situations? In a more general sense, when we speak of intelligence, we are referring to the ability to learn or understand from experience.

Thus, the answer to the question of *how* intelligent behavior comes about lies in understanding the nature of these experiences. There is no question that this "ability" to learn is inherited, that is, all healthy humans possess it fairly equally, and its physiological basis is the cerebral cortex. The component that has been utterly ignored by most AI researchers and cognitive scientists alike is the role of experience which is arguably the most important of these three necessary ingredients in learning. Especially for humans, experience is our teacher, it is the "programmer" of our brains, and it is the key to understanding our almost infinite behavioral adaptiveness and flexibility. It is thus the key to understanding what we usually mean by natural intelligence.

Now that we have a working definition of intelligence, it is important to correct some logical errors we make in our everyday language concerning intelligence. First of all, when we speak of intelligence, we are not speaking of a thing or an entity that we possess and that determines our behavior. The term is an abstraction. In other words, it is a word that we learn to use to refer to certain commonalities observed about behavior in its context. For example, on the basketball court we call Magic Johnson's behavior intelligent when we see him running down the court and passing the ball, without looking, to another player behind him. In the classroom, we describe as intelligent a student's behavior of correctly solving a problem. Likewise, if someone forgets the key to their house and uses a credit card to open the lock, we call that behavior intelligent. When a two-year-old child continues to look for an object that has been hidden from sight, we describe that as intelligent behavior. When a chess player finds the one move that places his or her opponent into checkmate, we call that behavior intelligent. And, when a human or a computer infers the correct chemical structure given a large amount of mass spectrographic data, we call that intelligent. These and all other examples of behavior in its context illustrate that the word intelligence is just that,

a word that we use to refer to certain commonalities found in behavior. Interestingly, it may be impossible to specify precisely these commonalities, that is, the precise stimulus boundaries determining whether we call some behavior intelligent or not. This has been supported by experimental evidence on the development of conceptual behavior in nonhumans. The linguistic point about our use of the word intelligence is an important one because it implies that there is no such thing as natural intelligence if by that we mean a quality that an organism possesses and that determines that organism's behavior. Hence, when we maintain that intelligence is an aptitude, or a "faculty of thought and reason," or "superior powers of mind" (Berndt, 1992), we are buying into the notion that intelligence is a real thing that we possess in varying quantities. It is possible that we have been fooled by our own language about ourselves when we believe that words such as intelligence refer to real things located inside us just as the word chair refers to a real tangible object. Moreover, we have also been fooled by our own language when we believe, as Plato suggested, that there exists in this universe some essential concept of chairness or intelligence. The only real tangible thing that the word intelligence refers to is overt behavior (i.e., intelligence is a tact), and the behaviors and contexts that cause us to utter the word "intelligent" are as varied as the number of different types of chairs that cause us to say "chair."

The implications of this simple change in our language are profound. For example, we can no longer say that intelligence tests measure intelligence; rather, they measure the correctness of behavior in a given, albeit contrived, context. And, moreover, it is not correct to say that some observed behavior occurred because the person was intelligent, or because they possessed intelligence. These types of explanations simply take the name or label given to a behavior and turn it around as the explanation of that very same behavior, in other words, the explanations become circular. The bottom line is that having observed behavior

in its context, and even naming it as intelligent behavior, we are still left without really understanding it. Understanding behavior in the scientific sense means being able to specify the conditions or variables necessary for its occurrence. This is what experimentation in science accomplishes. These specifiable conditions are not to be found in the mind (which itself cannot be found) or in the brain (which is neither structurally nor functionally like a digital computer), or even in rules or symbols (which are themselves abstractions). Rather, the conditions responsible for the occurrence of behavior are to be found in the experiences of the behaving organism. And there is only one discipline which defines as its subject matter the scientific study of the effects of experience on behavior, and that is the field of behavior analysis. For behavior theorists, experience means the known laws and principles that make up behavior theory. More specifically, experience means the contingencies of reinforcement in a particular environment that are responsible for adaptive behavior. And, although nonbehavior-analytic AI researchers have been slow to come around to this point, some have recently done so (e.g., Todd, 1992).

A SELECTIONIST MODEL OF NATURAL INTELLIGENCE (LEARNING)

A functional view not only of behavior and experience but of the language used to describe them has important implications for efforts to simulate human intelligence. For one, it means that any AI device must be an adaptive one, that is, it must learn. And, because the only demonstrated scientific laws of learning are found in what is known as Pavlovian and operant conditioning, the AI machine must somehow be programmed to learn in these ways. Specifically, a program must be able to learn through operant conditioning which is, after all, at the heart of behavior that we refer to as intelligent, purposeful, and intentional (Skinner, 1974). This means that a machine's "behavior" in a specific context must be sensitive to its own conse-

quences, that is, its effects on the environment. Much of what behavior theorists know about operant conditioning has been gleaned from the study of nonhuman animals, such as rats and pigeons, in controlled environments. Consider a simple example.¹ Suppose we put a hungry pigeon in an operant conditioning apparatus. On one wall above a food dispenser is a plexi-glass disc (called a key) that consists of two halves that can be illuminated either the same or different colors and that, when depressed, can operate a feeder. Now suppose we arrange it that such that when the two halves of the key are the same color, say either all red or all violet, then a peck on the key by the pigeon will operate the feeder. Conversely, when the two halves of the key are different colors, for example, one half red and the other half violet, pecking the key does not operate the feeder. Soon, the pigeon only pecks the key when both halves are the same color. Now let's test the pigeon with new colors such as blue-blue, yellow-yellow, blue-yellow, and yellow-blue. We will observe that the bird pecks the key more when the two halves are the same color than when they are different colors. Upon observing this outcome, most people would say that the pigeon is behaving intelligently. Initially we would say that the pigeon discriminates all red or all violet from a combination of the two, that is, same from different. During testing, we might say that the pigeon learned the concept of same-different with respect to colors. To be consistent with our working definition of intelligence, we would say that the pigeon now perceives the similarities and differences between the two halves of the key. When we put the hungry pigeon back into this situation at some later time and observe the behavior occurring again, we might also say that the pigeon remembers the concept. These ways of talking about the pigeon's behavior are consistent with our long philosophical tradition of putting the causes of behavior inside the organism, for example, in the mind, and then attributing

the behavior or changes in the behavior to the organism itself. But of course the explanation for the behavior is not to be found inside the pigeon — in its mind or brain. The explanation of the origin of the behavior is to be found in the contingencies of reinforcement that define the particular environment in which the behavior was conditioned. The behavior of pecking when the two halves of the key were the same color and not when they were different colors was selected by the reinforcers programmed by the experimenter and it can be de-selected by the experimenter by disassociating the behavior from the food and the colors. Although we may describe the pigeon's behavior as intelligent, we certainly can not give the pigeon credit for the behavior without understanding that the reinforcement contingencies produced it and any changes in those contingencies will change the behavior. And what about the discrimination? Discriminate, from the latin *disciminare*, means to divide or to distinguish. Clearly it was the reinforcement contingencies programmed by the experimenter that were responsible for dividing or distinguishing the same from different colors. The pigeon's pecking behavior, which was initially characterized by substantial variation was selected by the demands of the reinforcement contingencies and consequently came under the control of the chromatic properties of the key.

This is one of the simplest demonstrations of basic adaptive behavior and probably no computer program has yet been written that can come close to simulating it. Dare we ask why? In order for a computer program to simulate this intelligence it must be able to formalize what is going on, and in order to do that the programmer must be well versed in the theories of behavior. And behavior theorists probably don't yet know enough to program a computer to simulate this type of learning, that is, to take into account variations in the response, in the level of food deprivation, in the amount of food, and in subtle changes in the stimuli that define the behavioral context. Knowledge of computers and the ability to describe the final per-

¹This example is based on an experiment by Malott and Malott (1970).

formance in terms of a formal set of rules or symbols has not and will never produce a machine with the learning capacity of even the laboratory pigeon, much less a human, unless the programmers understand the nature of a functional, selectionist theory of behavior.

SUMMARY AND CONCLUSION

As has been said, researchers in AI simply borrowed the wrong model for intelligent behavior. Following from at least 2000 years of Western philosophies of human intelligence which rely on unseen, intangible, invented internal constructs that have been made responsible for everything humans do, researchers in AI assumed that intelligence is a thing that we possess that is responsible for intelligent behavior and, moreover, that it could be formalized as a set of rules that could be used to program computers thus simulating general human intelligence. And not only that, many AI researchers were arrogantly unrealistic about what they knew and what they could accomplish. For example, as early as 1957, Herbert Simon wrote,

It is not my aim to surprise or shock you....But the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until — in a visible future — the range of problems they can handle will be coextensive with the range to which the human mind has been applied. (as cited in Dreyfus, 1979, p. 81)

And in 1967, Simon wrote that, "machines will be capable, within 20 years, of doing any work that a man can do" (as cited in Dreyfus, 1979, p. 129). It has been more than 20 years. Are these predictions any truer today than they were when they were made?

But has AI really done so badly? Have AI researchers really failed to produce machines (i.e., programs) that behave intelligently? Recall that intelligence is simply a word that describes certain commonalities in behavior and its context. Based on this simple functional view of our language regarding intelligence, the word can be used to describe the behavior of many AI programs and, from this perspective, AI

has succeeded. But if AI researchers want to achieve their lofty goal of simulating general human intelligence, that is, the behavioral plasticity that distinguishes human behavior, then they will have to abandon their current working information-processing model of human behavior and intelligence and adopt a more functional model. (It should be added that the new trend toward neural networks whose connections are strengthened or weakened as a result of the effects of the output are a step in the right direction. In fact this demonstrates that the field of AI is itself affected by selectionist principles.)

CURRENT TRENDS IN AI: MOVING TOWARDS A SELECTIONIST MODEL

Behavior analysts might expect that adopting a functional model will lead AI workers to the only discipline that has achieved even a modicum of scientific understanding of human behavioral adaptiveness, the experimental analysis of behavior. However, the jury is not yet in. These and other points were discussed in a recent issue of *Computer* magazine, in a section titled "Hot Topics." For example, Conrad (1992) begins his contribution with the following sentence:

Computer technology has opened new approaches to many scientific disciplines, but the implications are more profound for the biological sciences because they lack the strong mathematical traditions of the physical sciences. (p. 79)

Then, in the process of describing the pursuits of his colleagues to construct a conceptual biological (neuronal) model algorithmically, Conrad describes three types of "learning." Behavior analysts might then expect a description of operant and respondent learning or, perhaps, observational learning. However, the types of learning Conrad describes are odd. In the first type of learning a neuronal repertoire is "created and tuned." According to Conrad "this can be achieved through a variation-selection evolution process, which acts on variations affecting neurons' internal dynamic parameters" (p. 79). In the second type of learning the neurons "must be selected and combined into use-

ful groupings, " which, according to Conrad, "requires a learning algorithm that operates on the access routes to these neurons." Conrad calls the third type of learning an "associative capability," which represents the ability of organisms to alter the linkage between perceptual experience and action in a rapid and flexible manner." Although it might not be possible to map precisely Conrad's three types of learning onto what behavioral psychologists have discovered about learning, there appears to be some overlap. The first two types of "learning" seem more analogous to the selection of genes through natural selection. However, the third type of learning is closest to what behavior theorists speak of as operant conditioning. The "ability" of organisms Conrad speaks of is what behavior analysts refer to as the capacity for behavior to be operantly conditioned. Of course, it is not the organism that "alters the linkage between perceptual experience and action;" it is the selective action of the environment on behavioral variation.

In the same issue of *Computer*, however, Todd (1992) describes an approach to machine intelligence that more closely maps onto behavior theory and operant conditioning. For example, Todd (re)defines intelligence as,

adaptive behavior: achieving the goals and satisfying the desires and motivations of the individual animat. The goals in turn reflect the actions necessary to survive and succeed in the animat's environment. (p. 78)

Elsewhere, Todd (1992) describes the importance of considering learning in designing AI machines:

If creatures or agents can pick up key features of their environment and figure out what to do about them, that is, learn how to behave adaptively, then much less of their behavioral repertoire must be built in or "hardwired," either by hand or through an evolutionary process. (p. 81)

Of course, as has already been stated, creatures don't "figure out" what to do about the environment, the environment selects from a range of behaviors. Notwithstanding this small semantic point, Todd (1992) is on the right track, moving away from an

essentialist conception of natural and artificial intelligence toward a selectionist one.

Of course, behavior analysts need not necessarily wait for AI researchers to "come around" to a behavior analytic philosophy of intelligence and learning; they can take the initiative and begin applying behavior-analytic theory to the problems at hand. The two papers that follow represent an important step in that direction, with a specific emphasis on the problem of programming computers capable of acquiring verbal relations. Cherpas (1992) suggests that truly effective natural language processing systems won't be possible unless researchers adopt a general selectionist approach to the behavior of the individual and, in particular, use Skinner's functional analysis of language presented in his book *Verbal Behavior* (1957). Stephens and Hutchison (1992) describe a specific attempt to construct a neural (they call it adaptive) network software system based on behavior-analytic principles that can, with the collaboration between commercial and academic individuals within the behavioral community, "be extended to provide a new generation of software systems capable of acquiring verbal behavior" (p. 151).

By thus adopting a selectionist, biological, behavior-analytic model of behavior instead of an essentialist, cognitive, information-processing model, AI researchers might be able to make some progress toward their lofty goal of simulating what we might refer to as general human intelligence.

REFERENCES

- Berndt, T. J. (1992). *Child development*. Ft. Worth: Harcourt, Brace and Jovanovich.
- Cherpas, C. (1992). Natural language processing, pragmatics, and verbal behavior. *The Analysis of Verbal Behavior*, 10, 135-147.
- Conrad, M. (1992). Toward an artificial brain. *Computer*, 25, 79-80.
- Crossman, E. K. (1985). The kiss and the promise: A review of Hubert L. Dreyfus' *What computers can't do: The limits of artificial intelligence*. *Journal of The Experimental Analysis of Behavior*, 44, 271-277.
- Dreyfus, H. L. (1979). *What computers can't do: The limits of artificial intelligence*. (Rev ed.) New York: Harper Colophon.
- Futuyma, D. J. (1979). *Evolutionary biology*. Sunderland, MA: Sinauer.

- Gardner, H. (1985). *The mind's new science: A history of the cognitive revolution*. New York: Basic Books.
- Malott, R. W., & Malott, M. K. (1970). Perception and stimulus generalization. In W. C. Stebbins (Ed.), *Animal psychophysics: The design and conduct of sensory experiments*. New York: Appleton-Century-Crofts.
- Palmer, D. C., & Donahoe, J. W. C. (1992). Essentialism and selectionism in cognitive science and behavior analysis. *American Psychologist*, 47, 1344-1358.
- Sack, J. M. (1984). Cognitive science research. In T. O'Shea and M. Eisenstadt (Eds.), *Artificial intelligence: Tools, techniques, and applications* (pp. 155-177). New York: Harper & Row.
- Searle, J. (1984). *Minds, brain and science*. Cambridge, MA: Harvard University Press.
- Skinner, B. F. (1957). *Verbal behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Skinner, B. F. (1974). *About behaviorism*. New York: Alfred Knopf.
- Stephens, K. R., & Hutchison, W. R. (1992). Behavioral personal digital assistants: The seventh generation of computing. *The Analysis of Verbal Behavior*, 10, 149-156.
- Todd, P. M. (1992). Two approaches to machine intelligence: The animat path to intelligent adaptive behavior. *Computer*, 25, 78-81.