non-monotonic reasoning. *Artificial Intelligence*, *13*, 27–39.

McCloskey, M., & Cohen, N.J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychological Learning and Motivation: Advances in Research and Theory*, *24*, 109–165.

Minsky, M. (1986). *The society of minds*. New York: Simon and Schuster.

Newell, A. (1973). Production systems: Models of control structures. In W.C. Chase (Ed.), *Visual information processing* (pp. 463–526). New York: Academic Press.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Newell, A., Shaw, J.C., & Simon, H.A. (1958). Elements of a theory of human problem solving. *Psychological Review*, *65*, 151–166.

Newell, A., & Simon, H.A. (1963). GPS, a program that simulates human thought. In E.A. Feigenbaum & J. Feldman (Eds.), *Computers and thought* (pp. 279–293). New York: McGraw-Hill.

Nilsson, N. (1980). *Principles of artificial intelligence*. Palo Alto, CA: Tioga Press.

Palmer, S. (1978). Aspects of representation. In E. Rosch & B.B. Lloyd (Eds.), *Computing and categorization* (pp. 259–303). Hillsdale, NJ: Erlbaum.

Powell, D.J., Tong, S.S., & Skolnick, M.M. (1989). EnGENEous domain independent, machine learning for design optimization. *Proceedings of the Third International Conference on Genetic Algorithms*. Los Altos, CA: Morgan Kaufman.

Powers, W.T. (1973). *Behavior: The control of perception*. Chicago: Aldine.

Powers, W.T. (1978). Qualitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, *85*, 417–435.

Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, *98*, 285–308.

Rosenbloom, P.S. (1983). The chunking of goal hierarchies: A model of practice and stimulus-response compatibility. Pittsburgh, PA: Ph.D. dissertation, Computer Science Department, Carnegie-Mellon University.

Rumelhart, D.E., & McClelland, J.L. (1986). On learning the past tenses of English verbs. In D.E. Rumelhart, J.L. McClelland, and the PDP research group, *Parallel distributed processing: Explorations in the microstructure of cognition*, *2*, 216–271.

Rumelhart, D.E., McClelland, J.L., and the PDP research group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* (3 vols.), Cambridge, MA: MIT Press.

Sejnowski, T.J., & Rosenberg, C.R. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, *1*, 145–168.

Simon, H.A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, *26*, 467–482.

Simon, H.A. (1990). *Models of my life*. Cambridge, MA: Harvard University Press.

Smolensky, P. (1987a). On variable binding and the representation of symbolic structures in connectionist systems. Technical Report CU-CS-355-87, Department of Computer Science, University of Colorado at Boulder.

Smolensky, P. (1987b). The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn. *Southern Journal of Philosophy*, *26*, 137–163. (Special issue on connectionism and the foundations of cognitive science).

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, *11*, 1–74.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, *46*, 159–216.

Touretzky, D.S. (1990). BoltzCONS: Dynamic symbol structures in a connectionist network. *Artificial Intelligence*, *46*, 5–46.

Waldrop, M.M. (1989). Humanity 2, computers 0. *Science*, *246*, 572–573.

Waltz, D.L., & Pollack, J.B. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science 9*, 51–74.

Commentary to Feature Review

PSYCHOLOGY AS A COGNITIVE SCIENCE

Dominic W. Massaro

University of California, Santa Cruz

My concern is with the role of psychology within cognitive science and how psychology has been changed by cognitive science. Each discipline within cognitive science has a different agenda, and the wholesale borrowing of models from one domain to another is not necessarily productive. Testing the psychological reality of normative, linguistic, logical, or economic models is limited because these models must first be modified to include psychological constraints. A more profitable tack is to develop and test descriptive models based on human performance and also to analyze the normative properties of these models. I see a continuity of psychological inquiry throughout the last century, and a continuing need for a functional level of behavioral description. This is the goal sought by much of psychology, and one that can be achieved perhaps with the help of other disciplines. To date, there is no reason to believe that the role of psychology and the need for a functional description will diminish.

When the editor asked me to contribute my thoughts on the unity of cognitive science, I agreed without hesitation. What could be easier? What I discovered is that doing science is easier than reflecting on the science one is doing. My goal is to give some observations on the discipline of psychology, how it has been influenced by other sciences, and how it might exist within cognitive science. Experimental psychology focuses on the method of inquiry, whereas cognitive psychology stresses the content of the inquiry. Cognitive science might be considered a blending of these two descriptions because its name implies both method and content. Of course, cognitive science also consists of several disciplines. One question for us is whether psychology has become more

Address correspondence to Dominic W. Massaro, Program in Experimental Psychology, University of California, Santa Cruz, CA 95064.

interdisciplinary because of its alignment with cognitive science, or whether it has retained its disciplinary framework.

A LINK WITH THE PAST: THE CONTINUITY OF INQUIRY

William James defined psychology as the science of mental life, whereas John B. Watson defined it as the science of observable behavior. The twentieth century has revealed that these two definitions are not necessarily incompatible with one another. What our field is called is not as important as what we as psychologists do. I have witnessed the description of our field change from experimental psychology to cognitive psychology to cognitive science. Soon after Neisser's (1967) book, *Cognitive Psychology*, obtained its early recognition and impact, George Sperling announced that he was now being called a cognitive psychologist even though his research hadn't seemed to change. We have at least one documentation (Lovie, 1983) that attention (a cognitive concept) received as much research during the heyday of behaviorism as it did during the heady days of cognitive psychology. Neisser (personal communication, March 1991) admitted that his doubts about the veracity of the cognitive revolution promised by his book were strengthened when the behaviorists themselves saw no great discontinuity between this framework and behaviorist research and theory (Salzinger, 1973). What these anecdotes reveal to me is the continuity of psychological inquiry throughout the last century. When faced with a doughnut, either the hole or the dough can be viewed as figure. Even so, I will argue that the goals of our discipline have remained fairly stable while the content and methods of inquiry have been refined and broadened in a cumulatively productive manner.

My argument for the continuity in psychological inquiry might cast Neisser's (1976) later book, *Cognition and Reality*, as the revolutionary one because he was promoting the overthrow of the experimental paradigm that had been refined from the procedures of Donders, Fechner, and Ebbinghaus. In fact, Neisser criticized the methods of experimental psychology because most experiments tended to ignore important questions. Neisser's rejection of most laboratory research because it was not "ecologically valid" challenged some of the central assumptions of experimental psychology. The debate initiated by this call is still active (Banaji & Crowder, 1989; Loftus, 1991; Neisser & Winograd, 1988; Roediger, 1990), although it should be clear to all that we are seeking laws to describe regularities in nature, not just ecologically valid experiments. The fact that these laws might only be revealed in well-controlled, artificial laboratory situations does not make them irrelevant to explanations of everyday life. The laws will help us describe, simulate, and understand complex natural phenomena, even though the laws will not necessarily be directly testable or even observable in the domain of those same phenomena (Manicas & Secord, 1983).

After the seventies decade of cognitive psychology, the eighties bowed to cognitive science. No one will argue with the similarities between these two, although the latter is necessarily more interdisciplinary than the former. Psychology has always been influenced by other disciplines, but perhaps the other disciplines were less influenced by psychological inquiry. A formalization of the cognitive sciences has made cognitive psychology more familiar to scientists from other disciplines and, therefore, seems to have placed psychology in a more influential position. In artificial intelligence, for example, more lip service is now paid to how artificial devices might be improved by considering how humans perform the behaviors of interest. In speech recognition by machine, devices are being developed to read the lips of the talker because humans use this information in face-to-face communication (Petajan, 1985).

Analogous to the continuity between the behaviorist and cognitivist eras, we should not view cognitive science as a paradigm shift for cognitive psychology. It appears to me that the scientists working within cognitive science have maintained their home-disciplinarian status, even though they are now part of a larger conglomerate of cognitive science. Contrary to Gardner's (1985, p. 390) vision, scientists within a discipline have common goals that cut across their specialty content areas. Thus, a psychologist studying language and another studying perception should have more common ground than the first psychologist and a linguist or the second psychologist and a philosopher of epistemology. Furthermore, interdisciplinary cooperation is not original with the founding of cognitive science. In terms of notable researchers, psychology has always been interdisciplinary: Donders was an optometrist; Wundt, Fechner, and Ebbinghaus philosophers; Helmholtz a physicist; and Hering a biologist. If we enumerate some influences, the theory of signal detection was purchased wholesale from engineering, prototype theory pulled from Wittgenstein's (1953) philosophical treatise, and so on. Some exemplary cases of interdisciplinary studies of psychological questions are the classic papers by Selfridge and Neisser (1960) on pattern recognition, Miller and Chomsky (1963) on language, and Estes and Suppes (1959) on mathematical psychology.

DISCIPLINARY AGENDAS WITHIN COGNITIVE SCIENCE

If one simply scans the literature, it becomes obvious that the disciplines within cognitive science have different agendas (see Simon and Kaplan, 1989, for an informative analysis of the metatheories of different disciplines and their mutual influence). Notwithstanding the different goals, psychology has constantly adopted the theories of other disciplines as potential psychological models. The psychology of language provides an illuminating example. Linguistics desires an elegant grammar that generates all and only the permissible sentences of a language. Psychology should describe and explain how language is used. An important interdisciplinary event, perhaps, was the founding and blossoming of the field of psycholinguistics. However, psycholinguistic inquiry was bogged down by testing the psychological reality of linguistic theory rather than concentrating on developing and testing hypotheses about how language is actually used. Linguists relied on themselves for their database, whereas psychologists relied on linguists for their theories.

Although much of the work in the field of psycholinguistics is still wedded to the testing of linguistic theory, it is much healthier today. An example of a productive approach is the functional grammar framework (MacWhinney & Bates, 1989). The theoretical framework and experiments depend on both psychological and linguistic constraints, and they make testable predictions about both language use and acquisition. Another promising example of psychologically driven inquiry is Gernsbacher's (1990) uncovering of the processes involved in sentence comprehension. The Osherson and Lasnik (1990) volume, on the other hand, tends to maintain the tradition of formal linguistic theory and the testing of the psychological reality of that theory. The book is also representative of only the "modularity" perspective within linguistics and psycholinguistics. Thus, the chapters focus on the modularity of processes such as speech perception or simply articulate linguistic theory without even a concern for psychological reality of the theory. Needless to say, my bias tends to lean in the direction of homespun theories of language performance (Massaro, 1989b).

Not only linguistics but also philosophy, economics, and, to some extent, computer sciences have an arsenal of prescriptive or normative models like the propositional calculus, the rational man, and the Turing machine. Psychologists have also been concerned with testing the psychological reality of these models. Similar to the enterprise in linguistics, however, testing the psychological reality of these normative models might not be as productive as developing psychological models based on empirical results, and then asking to what extent these models might be considered normative. As an example, some investigators have documented situations in which people do not use base rates in decision making (Kahneman & Tversky, 1972), and have argued that decision making is not normative (optimal). Other investigators, on the other hand, have addressed the question of the psychological processes involved in decision making and how base-rate information is integrated with other sources of information in the decision-making task (Leon & Anderson, 1974).

Psychologists have also been influenced by philosophy and computer science for metatheory. The computer metaphor provided psychologists a working solution to the mind/body problem, one that did not deny mental events. Behaviorists had the reflex arc; cognitive psychologists had the sequential programming operations of the von Neumann computer. At last, a way to justify an analysis of the operations of the mind: They are analogous to the algorithms of a physical computational device. For historical correctness, however, many psychologists, with their nose to this information-processing grindstone, did not buy into computation in the Turing machine sense. Although psychologists were using the computer as a tool, its role as a theory did not naturally follow the reasonable proposition of "tools into theories" (Gigerenzer, 1991).

Palmer and Kimchi (1986) sketch five properties of the psychologist's information-processing framework, which illustrate that the psychological metatheory of information processing differs significantly from the notion of Turing computation within computer sciences. There is other evidence that much of the research in cognitive psychology was not particularly tied to the formal notion of computation (Heijden & Stebbins, 1990). Theories and studies of memory seem to be more influenced by the idea of spatial storage and search than on the computer metaphor directly (Roediger, 1980). In one case, at least, the sixties idea of feature analysis for pattern recognition was as influenced by observations in neurophysiology and linguistics as it was by pattern recognition by machines (Massaro, 1989a, pp. 315–318). Concepts such as familiarity, memory strength, stimulus-response compatibility—certainly not concepts derivative of the computer metaphor—continued to carry some theoretical burden during this time.

Other theorists were more directly influenced by the architectures and algorithms being used in computer science and artificial intelligence. Some psychologists adapted production systems as their currency for information-processing theory. Because any theory of this form has the potential of being a Turing equivalent (Newell, 1990), the theorist's goal is to provide as many constraints on the theory as possible. Newell's (1990) SOAR is the most impressive effort to date within this genre. SOAR might be viewed as a formalization in production systems of many of the psychological principles uncovered in the last decades of research. It also maintains the sequential stage structure inherent in the information-processing approach. To enhance its impact, I would advise Newell and company to do what the PDPers accomplished with their third book and the two little floppies in the back (McClelland & Rumelhart, 1988). Such easy access to the systems might open the door for testing and enhance their acceptance as models worthy of consideration.

Psychology has now witnessed a transition from the computer metaphor of information processing to the brain metaphor of neural processing, with the rebirth of connectionism. (However, one might point to the connection machine [Hillis, 1987] as the metaphor in which case connectionism simply substituted one type of computer metaphor for another.) This idea converted both the innocent and those who should have known better. The promises of brain-style modeling multiply with every new conversion. One would like to believe that current brain-style modeling has been a significant advance over more traditional functional models. Alas, this seems not to be the case. First, the networks and their "learning" algorithms are not very brainlike and, in fact, no more brainlike than production systems (see also Anderson, 1990; Mewhort, 1990). This neural fantasy became obvious enough so that Smolensky's (1988) sellout of neural modeling for subsymbolic modeling was the only viable route for the new connectionism. Second, a network's solution to a problem is called learning, but it might be interpreted more naturally as a minimization routine that attempts to optimize a mapping between input and output. Third, a solution is virtually guaranteed if the theorist uses a sufficient number of hidden units and the operational system does not get trapped in a local minimum (Hornik, Stinchcombe, & White, 1989; Massaro, 1988, 1990). Fourth, a true emergence of knowledge in subsymbolic connectionism is as scarce as it is in production systems. Like multidimensional scaling, you get out what you put in. In the paradigm case, NETtalk "learned" to pronounce English text and its creators went on to discover that the hidden units had come to represent consonants and vowels as different classes (Rosenberg, 1987; Sejnowski & Rosenberg, 1986). However, this distinction is exactly what was programmed into the system by definition of its output units (Verschure, 1990).

Two positive spinoffs from neural networks offer hope for the future. First, connectionist modeling has restimulated psychologists' interest in learning, an area in which the potential for emergence of knowledge is at least possible, Mewhort (1990) also gives credit to the connectionists for the rebirth of the study of learning, which faltered during the cognitive revolution. Second, grounding our models in neurobiology remains a healthy enterprise as long as we are not easily fooled about what is fact and what is fiction (Crick, 1989). To improve the neurological plausibility of back propagation, for instance, Zipser and Rumelhart (1990) designed reciprocal cells to permit the two-way passage of information. Unfortunately, no neurological evidence for their new dual network is documented, and it appears to be an implausible structure.

PROXIMATE VERSUS ULTIMATE CAUSATION

Within the ideal world of the unity of science, psychology is about biological entities and evolutionary biology should contribute to our inquiry. In fact, it appears that psychology and evolutionary biology necessarily ask different questions. An important distinction made in evolutionary biology is between proximate and ultimate causes of (or influences on) behavior (Alcock, 1989). Proximate causes address the immediate environmental stimuli and the immediate processes that influence behavior. Ultimate causes address the evolutionary history of the organism. In simpler terms, proximate descriptions focus on *how*, whereas ultimate descriptions focus on *why*. For example, we might ask what environmental information the gannet (a large seabird) uses to signal closing its wings before landing on the water. Ultimate causes might address why the gannet closes its wings when landing—what evolutionary significance it might have. As psychologists, we have usually been concerned with proximate (immediate or close in time) influences. For example, what are the visual features actually used in letter recognition, how are these features combined, and how is a decision made given this information? Ultimate causes, such as the evolution of the visual system to detect edges and other properties of the letters, are of less interest. As illustrated by the above two approaches to the study of the visual system, the experimental psychologist's concern is with proximate causes. This concern with immediate causation, in many respects, makes the framework of evolution less applicable to psychological inquiry (Massaro & Friedman, 1991).

An interesting example of the importance of the distinction between proximal and ultimate causation has to do with the multiple cues and systems that are available for seeing the world in depth. Stereopsis developed rather late in evolution of the binocular (visual) system, perhaps because it evolved only to overcome camouflage. This process occurs rather early during the time course of processing a given scene, however, even though it occurred later in evolutionary time. Thus, evolutionary history (ultimate causation) and immediate causation (proximal causation) address different aspects of the phenomenon of seeing.

Convergent and divergent evolution illustrate that ultimate causation does not necessarily constrain proximal causation. Convergent evolution involves the independent acquisition over time through natural selection of similar characteristics in unrelated species. Divergent evolution involves the acquisition of dissimilar characteristics in the same species. Take as an example the act of seeing. As an example of convergent evolution, an octopus (a cephalopod) and a mammal evolved highly similar solutions to the problem of seeing (Blakemore, 1977). As a case of divergent evolution, some gull species nest on the ground in open grassy areas. When the breeding adults detect a potential egg-eating predator, they emit a volley of loud cries, fly toward it, divebomb, and defecate on the enemy. As an exception to this mobbing behavior, the kittiwake gull nests on nearly vertical cliffs on the coast. This site precludes mammalian predators and the erratic sea winds limit the threat of large predatory birds (Alcock, 1989). The kittiwake gulls do not mob the occasional enemy who drifts by the nesting site. Once again, the psychologist must explain the proximal causes of behavior whether it results from convergent or divergent evolution.

PSYCHOLOGY AT THE FUNCTIONAL LEVEL

My persistent requirement for a functional descriptive level of behavior reflects my belief that a part of our understanding is necessarily tied to proximal causation. Furthermore, the functional analysis made possible by the descriptive level will have implications for practice. As an example, consider an information-processing analysis of good and poor readers (Stanovich, 1986). Reading a sentence can be described as a sequence of processing stages going from transduction of the visible text to activation of various letter and word codes to arriving at some meaning of the utterance. It turns out that many poor readers tend to be deficient at the activation of letter and word codes (Perfetti, 1985), and one observes a larger influence of context for these readers. The corresponding events in the brain would be continuous, of course, and the brain level of description would differ from the information-processing one. Given the complexity of events at the brain level, some have argued against the possibility of a complete reductionist description (Massaro, 1989a; Uttal, 1990). Even if the brain-level description is possible, the information-processing analysis can provide an assessment and a diagnosis for remediation.

The information-processing level of analysis is appropriate for enhancing the quality of our everyday life, as witnessed by the substantial number of psychologists who have followed George Miller's advice and given psychology to the people (Adams, 1990; Greenwald, Spangenberg, Pratkanis, & Eskenazi, 1991; Norman, 1988). For the reader who has persisted during my prolonged comparison of psychology with other disciplines, we can ask where we have been and where we are going. Notwithstanding the numerous advances in many of the cognitive sciences, we expect that psychology will continue to develop its laws and regularities slowly but cumulatively. One important difference today is the existence of excellent undergraduate and graduate programs in cognitive science. Our students from these programs should be more receptive to and better prepared for interdisciplinary adventures. True interdisciplinary endeavor, however, will require significant compromises concerning the questions asked by each of the participating disciplines.

A LINK WITH THE FUTURE

Will psychology be consumed by cognitive science or some other discipline such as evolutionary theory? I don't think so, although I am sure psychology will continue to be influenced by developments in other fields. The influence from other disciplines must be critically assessed because the specific questions of one discipline differ significantly from those of another and, therefore, the solutions acceptable in the home discipline will rarely be sufficient for some other. In addition, we need to distinguish perspiration from inspiration. I have been discussing perspiration, whereas inspiration should have fewer bounds than the former. Why limit this influence and our horizons to the 5 + 2 cognitive sciences that we know so well? I take my inspiration where I can find it—whether it be Cézanne's reconstruction of the visual world, Dostoyevsky's stream of consciousness, a visual motif of a Mandelbrot set, Hawking's strong inference tests between creationist and the physical theories of the origin of the universe, or a sociobiologist's account of dimorphism. Like the ambiguous boundaries among sciences, let us acknowledge the similarities between the sciences and the arts/humanities.

The question of the unity of cognitive science versus a collection of unrelated disciplines parallels a distinction in categorization research: the one between summary description and exemplar theories of category representation. Our mental representation of a category might be described in terms of a summary description of a set of features. An instance belongs to the category to the extent that they share several features. According to exemplar theory, on the other hand, an instance is considered to be an instance of the category to the extent that the instance is similar to one or more exemplars of the category representation. Both the cognitive-science question and the categorization questions might remain unanswered for some time to come. With respect to cognitive science, however, perhaps an exemplar description (or at least a summary description composed of the disjunction of multiple descriptions) would be the most productive. Uniquely different disciplines would permit the assumptions and findings of one discipline to be evaluated in the framework of another discipline. A single discipline of cognitive science would preclude these multiple evaluations, making it more likely that the framework is not sufficiently questioned and tested. Cognitive science has magnified the telescopes of the disciplines that make it up and its unity may thrive on its dynamic disunity.

Acknowledgments—The writing of this paper was supported, in part, by grants from the Public Health Service (PHS R01 NS 20314) and the National Science Foundation (BNS 8812728). I thank Paul Verschure and Ray Gibbs for reading an earlier version of the paper and for helpful discussions.

REFERENCES

Adams, M.J. (1990). *Beginning to read.* Cambridge, MA: MIT Press.

Alcock, J. (1989). *Animal behavior.* Sunderland, MA: Sinauer.

Anderson, J.R. (1990). *The adaptive control of thought.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Banaji, M.R., & Crowder, R.G. (1989). The bankruptcy of everyday memory. *American Psychologist, 44,* 1185–1193.

Blakemore, C. (1977). *Mechanics of the mind.* New York: Cambridge University Press.

Crick, F. (1989). The recent excitement about neural networks. *Science, 337,* 129–132.

Estes, W.E., & Suppes, P. (1959). Foundations of linear models. In R.R. Bush & W.K. Estes (Eds.), *Studies in mathematical learning theory* (pp. 137–179).

Gardner, H. (1985). *The mind's new science.* New York: Basic Books.

Gernsbacher, M.A. (1990). *Language comprehension as structure building.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review, 98,* 254–267.

Greenwald, A.G., Spangenberg, E.R., Pratkanis, A.R., & Eskenazi, J. (1991). Double-blind tests of subliminal self-help audiotapes. *Psychological Science, 2,* 119–122.

Heijden, A.H.C. v.d., & Stebbins, S. (1990). The information-processing approach. *Psychological Research, 52,* 197–206.

Hillis, W.D. (1987). The connection machine. *Scientific American, 256 (June),* 108–115.

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks, 2,* 359–366.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology, 3,* 430–454.

Leon, M., & Anderson, N.H. (1974). A ratio rule from integration theory applied to inference judgments. *Journal of Experimental Psychology, 102,* 27–36.

Loftus, E.F. (1991). The glitter of everyday memory . . . and the gold. *American Psychologist, 46,* 16–18.

Lovie, A.D. (1983). Attention and behaviorism—fact and fiction. *British Journal of Psychology, 74,* 301–310.

MacWhinney, B., & Bates, E. (1989). *The crosslinguistic study of sentence processing.* New York: Cambridge University Press.

Manicas, P.T., & Secord, P.F. (1983). Implications for psychology of the new philosophy of science. *American Psychologist, 38,* 399–413.

Massaro, D.W. (1988). Some criticisms of connectionist models of human performance. *Journal of Memory and Language, 22,* 213–234.

Massaro, D.W. (1989a). *Experimental psychology: An information processing approach.* San Diego: Harcourt Brace Jovanovich.

Massaro, D.W. (1989b). Multiple book review of *Speech perception by ear and eye: A paradigm for psychological inquiry. Behavioral and Brain Sciences, 12,* 741–794.

Massaro, D.W. (1990). The psychology of connectionism. *Behavioral and Brain Sciences, 13,* 403–406.

Massaro, D.W., & Friedman, D. (1990). Models of integration given multiple sources of information. *Psychological Review, 97,* 225–252.

Massaro, D.W., & Friedman, D. (1991). Book review of *The adaptive character of thought* by John R. Anderson. *American Journal of Psychology, 104,* in press.

McClelland, J.L., & Rumelhart, D.E. (1988). *Explorations in parallel distributed processing: A handbook of models, programs, and exercises.* Cambridge, MA: MIT Press.

Mewhort, D.J.K. (1990). Alice in wonderland, or psychology among the information sciences. *Psychological Research, 52,* 158–162.

Miller, G.A., & Chomsky, N. (1963). Finitary models of language users. In R.D. Luce, R.R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (pp. 419–491). New York: Wiley.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

Neisser, U. (1976). *Cognition and reality.* San Francisco: W.H. Freeman.

Neisser, U., & Winograd, E. (Eds.). (1988). *Remembering reconsidered: Ecological and traditional approaches to the study of memory.* Cambridge, UK: Cambridge University Press.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Norman, D.A. (1988). *The psychology of everyday things.* New York: Basic Books.

Osherson, D.N., & Lasnik, H. (1990). *An invitation to cognitive science: Language,* Vol. 1. Cambridge, MA: MIT Press.

Palmer, S.E., & Kimchi, R. (1986). The information processing approach to cognition. In T.J. Knapp & L.C. Robertson (Eds.), *Approaches to cognition: Contrasts and controversies* (pp. 37–77). Hillsdale, NJ: Lawrence Erlbaum Associates.

Perfetti, C.A. (1985). *Reading ability.* New York: Oxford University Press.

Petajan, E.D. (1985). Automatic lipreading to enhance speech recognition. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, June, 19–23,* 40–47.

Roediger, H.L. III (1980). Memory metaphors in cognitive psychology. *Memory & Cognition, 8,* 231–246.

Roediger, H.L. III (1990). Book review of *Remembering reconsidered: Ecological and traditional approaches to the study of memory.* (Ed. by U. Neisser & E. Winograd) *American Journal of Psychology, 103,* 403–409.

Rosenberg, C.R. (1987). Revealing the structure of NETtalk's internal representations. *Proceedings of the ninth annual conference of the Cognitive Science Society* (pp. 537–554). Hillsdale, NJ: Lawrence Erlbaum Associates.

Salzinger, K. (1973). Inside the black box, with apologies to Pandora. A review of Ulrick Neisser's *Cognitive psychology. Journal of the Experimental Analysis of Behavior, 19,* 369–378.

Sejnowski, T.J., & Rosenberg, C.R. (1986). *NETtalk: A parallel network that learns to read aloud.* The Johns Hopkins University Electrical Engineering and Computer Sciences Technical Report. JHU/EECS-86-01.

Selfridge, O.G., & Neisser, U. (1960). Pattern recognition by machine. *Scientific American, 203 (August),* 60–68.

Simon, H.A., & Kaplan, C.A. (1989). Foundations of cognitive science. In M.L. Posner (Ed.), *Foundations of cognitive science.* Cambridge, MA: MIT Press.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences, 11,* 1–73.

Stanovich, K.E. (1986). Matthew effects in reading: Some consequences of individual differences in the acquisition of literacy. *Reading Research Quarterly, 21,* 360–407.

Uttal, W.R. (1990). On some two-way barriers between models and mechanisms. *Perception & Psychophysics, 48,* 188–203.

Verschure, P. (1990). *Connectionist explanation: Taking positions in the mind-brain tradeoff.* Unpublished paper.

Wittgenstein, L. (1953). *Philosophical investigations* (G.E.M. Anscombe, trans.). New York: Macmillan.

Zipser, D., & Rumelhart, D.E. (1990). The neurobiological significance of the new learning models. In E.L. Schwartz (Ed.), *Computational neuroscience* (pp. 192–200). Cambridge, MA: MIT Press.