CHAPTER 16

Cognitive science

PAUL THAGARD

There are many reasons why a budding academic might want to avoid interdisciplinary research. It is difficult enough to acquire expertise in one field of research, let alone two or more. The time required to read the literature in a field outside your own main area can be hard to find, and the additional time investment to learn novel methods from another field can be huge. Moreover, the hiring and reward systems in academia still run strongly along disciplinary lines, so that work that draws on or contributes to other fields may not be fully valued in one's own field. Interdisciplinary research may not be appreciated by narrow-minded colleagues. Some interdisciplinary projects have a bogus air about them, looking like they were designed more to bring in big research grants than to accomplish intellectual goals. The interdisciplinary scholar can look a bit like a dilettante, dabbling in multiple fields in order to avoid tackling the difficult problems in an established field. Grants for interdisciplinary research can be difficult to get, because most granting agencies are organized along disciplinary lines.

Despite these deterrents to interdisciplinary research, there are powerful intellectual reasons why work that oversteps the ossified boundaries of established fields can have great intellectual benefits. Such benefits are vividly apparent in the interdisciplinary field of cognitive science, which attempts to understand the mind by combining insights from the fields of psychology, philosophy, linguistics, neuroscience, anthropology, and artificial intelligence. After a brief review of the history of the field and its contributing disciplines, this chapter will examine some of the main theoretical and experimental advances that cognitive science has accomplished over the past half century, deriving lessons that might be useful for researchers in any emerging interdisciplinary area.

16.1 History

Construed broadly, cognitive science is as old as philosophical reflections about the nature of mind, and so dates back at least to Plato and Aristotle. Philosophers such as Francis Bacon, John Locke, David Hume, Immanuel Kant, and John Stuart Mill generated ideas about the contents and processes of thinking. Experimental psychology originated in the late nineteenth century with the establishment of laboratories by Wilhelm Wundt, William James, and others.

Modern cognitive science began in the 1940s when visionaries such as Alan Turing (1950), W. S. McCulloch (1965), Norbert Wiener (1961), and Donald Hebb (1949) began to apply emerging ideas about computing, engineering, and brain systems to develop new hypotheses about mental mechanisms. Previous mechanistic theories of mind, ranging from the atomism of Lucretius to the behaviorism of B. F. Skinner, were much too impoverished to explain the complexities of human thinking. But in the mid-1950s there emerged a panoply of powerful ideas about how mental processes could be understood by analogy to computational ones. The major contributors included the psychologist George Miller (1956), the linguist Noam Chomsky (1957), and researchers in the nascent field of artificial intelligence, including Herbert Simon, Allan Newell, Marvin Minsky, and John McCarthy (McCorduck 1979). The year 1956 was particularly notable, as it marked publication of Miller's famous paper on information processing, 'The magical number 7 plus or minus 2', and the Dartmouth Conference that initiated the field of artificial intelligence. The fundamental hypothesis of cognitive science, that thinking consists of computational procedures applied to mental representations, began to influence research in psychology and other fields.

The term 'cognitive science' was only coined two decades later (Bobrow and Collins 1975). Events in the late 1970s included the formation of the Cognitive Science Society, the creation of the journal *Cognitive Science*, and the establishment of cognitive science programs at many universities. Today, evidence that interdisciplinary research and teaching in cognitive science is thriving includes multiple successful journals, international societies with regular conferences, and active teaching and research programs in many universities and organizations around the world. For detailed treatments of the history of cognitive science see Gardner (1985), Thagard (1992, 2005b), and especially Boden (2006).

16.2 Patterns of collaboration

The interdisciplinary structure of cognitive science is displayed in the hexagon in Fig. 16.1, the original version of which appeared in a report for the Sloan Foundation in 1978 (Gardner 1985, p. 37). The 13 lines in the hexagon indicate the range of possible connections between the six main disciplines of cognitive science, but the links are misleading in several respects. First, the disciplines have been highly unequal participants in interdisciplinary research. For example, although anthropology has contributed some highly interesting work on mental representations and processes in non-Western cultures, most anthropologists have shown little interest in cognitive science. More significantly, some of the most widely read philosophical discussions of cognitive science have been highly critical of it, for example attacks by Herbert Dreyfus (1979) and John Searle (1980) on the computational view of minds. The field of artificial intelligence has moved away from the interest in human thinking that inspired its early decades to a more engineering-oriented concern with the building of intelligent computers. In contrast, most cognitive psychology research is naturally dedicated to understanding the operation of human intelligence.

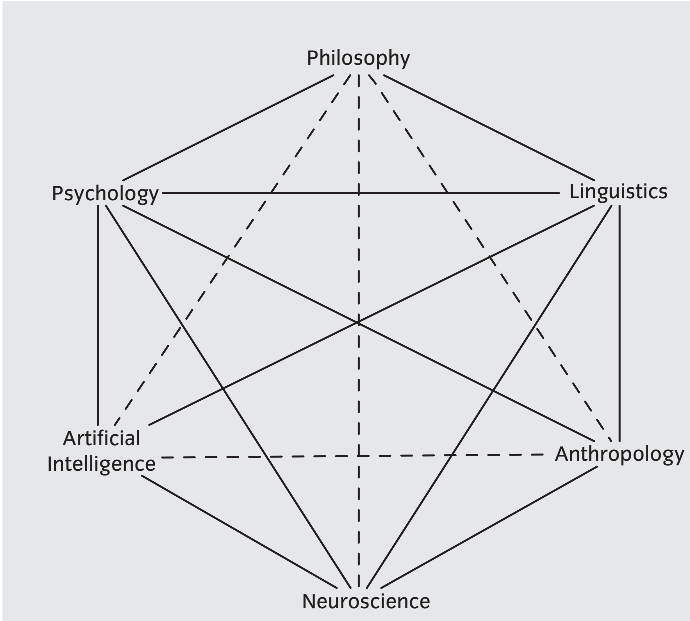

Figure 16.1 Connections among the cognitive sciences (based on Gardner 1985, p. 37). Unbroken lines indicated strong interdisciplinary ties *c*. 1978, and broken lines indicate weak ones. The ties between philosophy and both neuroscience and artificial intelligence are much stronger today.

Second, the hexagon does not convey the historical fact that some combinations of the fields have been much more active than others and that levels of activity have varied over time. When cognitive science began officially in the 1970s, by far the most prominent kind of interdisciplinary collaboration occurred at the intersection of psychology and artificial intelligence, continuing a pattern established in the 1950s by pioneers such as Herbert Simon (1991). Psycholinguistics also flourished early on. Neuroscience became much more central starting in the 1980s and 1990s, with the increased sophistication of neurally inspired computational models and the development of brain scanning technology that greatly expanded the possibilities for neuropsychological experiments. Philosophers' involvement in cognitive science has been highly variable, ranging from dismissal on the grounds that philosophy must transcend the merely empirical (Williamson 2007), to systematic reflection on controversial issues such as the extent to which knowledge is innate (Stainton 2006). Since the 1980s there has been much philosophical discussion of issues that arise in cognitive psychology and neuroscience (e.g. Thagard 2007; Bechtel 2008). Most strikingly, the application of psychology and neuroscience to traditional philosophical problems in ethics and epistemology has become an active enterprise (e.g. Appiah 2008; Knobe and Nichols 2008; Thagard 2010). For example, progress in neuroscience raises serious challenges to traditional ideas about free will and responsibility. In contrast, philosophers' interest in linguistics has waned, probably because language is no longer seen as so central to philosophy as it used to be; and work at the intersection of philosophy and anthropology has always been rare.

The third misleading feature of the hexagon is that the lines only indicate binary relations between disciplines, whereas some important developments have involved collaborations across several fields. For example, computational psycholinguistics draws on ideas from three disciplines to develop formal models of how minds use language. Current work in theoretical neuroscience combines the study of brains with psychological and computational ideas. Recent work on emotion attempts to address philosophical issues about rationality by means of computational models that are psychological, neurological, and even sometimes social (Thagard 2006). In sum, although Fig. 16.1 provides a useful diagram of possibilities for interdisciplinary connections, it does not display the shifting patterns of disciplinary involvement in such research.

There are at least three styles of interdisciplinary interconnection. The first is when an individual alone does research at the intersection of two or more disciplines. This requires the researcher to acquire mastery not only of the ideas but also of the methods of more than one field. For example, there are psychologists who have learned to do computational modeling, and a few philosophers who have learned to do experiments in psychology or neuroscience.

A second powerful kind of interdisciplinary interconnection involves collaboration, in which two or more individuals work together on a project combining their knowledge and skills in ways that require some mutual comprehension but not full duplication of abilities. This pattern of research has often been the most successful one in cognitive science, which has benefited from collaborations involving people whose original backgrounds combined, for example, psychology and artificial intelligence, psychology and neuroscience, and linguistics and anthropology.

The third style of interdisciplinary research does not require such collaboration or even individuals who have mastered more than one field. There has been much valuable work by more narrowly disciplinary researchers that draws on ideas from related fields. For example, Eleanor Rosch's influential work on concepts as prototypes was inspired in part by ideas of the philosopher Ludwig Wittgenstein (Rosch and Mervis 1975). Many articles published in the journal *Cognitive Science* are not internally interdisciplinary, as they lack a combination of methods. However, most articles that appear there are intended to be of interdisciplinary interest in that they address concerns inspired by or relevant to work in various fields concerned with the nature of mind and intelligence. For example, an experimental paper on the nature of human concepts falls squarely within cognitive psychology, but should be relevant to philosophical, computational, neurological, linguistic, and cross-cultural issues about mental representations. This third style of interdisciplinary research requires less personal investment than the individual mastery and collaborative styles, but it usually presupposes at least some acquaintance with relevant literature in other fields.

In the introduction, I mentioned some of the impediments to interdisciplinary research, but have described how cognitive science has provided a strong example of a successful effort to combine insights and methods from at least six disciplines. Now I want to depict more fully what that success has consisted in, by discussing the theoretical and experimental benefits of being interdisciplinary.

16.3 Theoretical benefits

A scholar has been defined as someone who knows more and more about less and less. Pursuing minutiae is often an effective strategy in academic research, since becoming an expert in some narrow niche is often a good way to publish and secure tenure. For the more intellectually ambitious, however, it is much more exciting to pursue theoretical ideas that are both important and novel. How can such creativity be achieved?

It helps, of course, to be a genius, with cognitive resources such as unusually powerful memory, imagery, or speed in connecting previously unrelated ideas or facts. But creativity is not only for the swift, because others of more modest intellectual capacities can still be creative by putting together ideas that have not been associated by other thinkers. Perhaps it takes a genius to work in a well-trodden area and manage to come up with something totally novel, but for the rest of us there is an easier road to creativity. Instead of focusing narrowly on one academic field, a researcher can cast a broader intellectual net and make new connections by tying together ideas from different disciplines. Cognitive science has thrived intellectually by making such creative theoretical connections.

In the mid 1950s, the dominant psychological theories, especially in the United States, were behaviorist, claiming that a scientific approach to the mind should restrict itself to considering how environmental stimuli are correlated with behavioral responses. Behaviorism was encountering difficulties in explaining the complex performance of rats, let alone humans, but theories are rarely rejected because of empirical problems alone. Rather, it is only when an alternative theory comes along with a new way of explaining recalcitrant data that a dominant theory comes strongly into question (Kuhn 1970; Thagard 1992). What happened around 1955 was that ideas from the rapidly emerging study of computers provided a new way to think about mental processes that was as rigorously mechanistic as behaviorism but possessed much more explanatory power.

A computer program consists of a set of structures, such as numbers, words, and lists, and a set of algorithms, which are mechanical procedures that operate on those structures. Those not familiar with computer programs can think of how people add up a list of numbers, where the structures are the numbers and the algorithm is the procedure for addition learned in elementary school. Or consider a recipe book, in which the recipe consists of a list of ingredients (the structures) to which people apply a set of procedures such as mixing and baking. Computer programs provide a highly suggestive analogy about how minds might work: mental representations may be like the structures used in computer programs, and mental procedures may be like the algorithms that make computers run. The strongest claim to consider is not only that thinking is like computing, but that thinking in fact is a kind of computing (Thagard 2005b).

The analogy just described has been fertile in suggesting many new ideas about how representational structures and computational procedures might be responsible for mental processes such as perception, memory, learning, problem solving, language use, and so on. Many productive specific theories have been developed about how rules, concepts, images, and analogies might operate in the mind. This theoretical productivity could never have happened if psychologists had stuck with the intellectual resources of behaviorism. Instead, by importing ideas from the study of computers, it became possible to formulate creative new theories of mental functioning. Whereas behaviorism restricted itself to stimulus–responses connections, cognitive science investigates how behavior and thought result from mental representations and computational procedures that integrate perceptual stimuli and produce responses based on complex inferences.

Another interdisciplinary source of ideas about how the mind works is the study of the brain. Some early ideas about how the mind works drew on neural mechanisms, but brain-style computing only took off in the 1980s through the development of an approach known as connectionism or parallel distributed processing. Brains operate differently from conventional computers. Neurons are slow, firing on average fewer than 100 times per second, but they perform powerful computations by virtue of the fact that that there are so many of them (around 100 billion) operating in parallel. In contrast, computer chips are very fast, with billions of cycles per second, but they usually operate serially, one step at a time. Today there is a flourishing field called theoretical neuroscience that develops new computational ideas about how brains support various kinds of thinking (Dayan and Abbott 2001).

Besides computer science and neuroscience, psychology has also been influenced by ideas from other fields, including philosophy and linguistics. Psychology is not just a recipient of theoretical ideas, but has also served as a donor. Psychology has contributed to the field of artificial intelligence that tries to build computers capable of some of the impressive feats of problem solving accomplished by people. For example, some expert systems that are engineering projects with the aim of making computers capable of tasks such as medical diagnosis have drawn on psychological ideas about mental representations like rules, analogies, and neural networks. Philosophy of mind and cognitive anthropology have also been heavily influenced by developments in cognitive psychology. Oddly, cognitive science has had little influence on fields such as literary theory and history, which could greatly benefit from richer ideas about how minds find meaning and make decisions.

Many more specific examples of the development of new theoretical ideas in cognitive science through interdisciplinary collaboration could be given, but here are two illustrations. The study of analogy has blossomed since the 1980s as the result of theoretical ideas that have combined insights from philosophy, psychology, artificial intelligence, and neuroscience. The goal of trying to understand how minds can often so productively apply ideas from one domain to another was studied by philosophers such as Mary Hesse (1966), but was greatly fostered by the development of new psychological ideas about how minds can use representations of one problem to solve another. Psychologists such as Dedre Gentner and Keith Holyoak devised new ideas about how people use analogies, partly on the basis of their own experiments but also drawing heavily on computer models, including ones that employ artificial neural networks (e.g. Gentner 1983; Holyoak and Thagard 1995; Gentner *et al.* 2001).

Recent work on emotion has also been highly interdisciplinary, drawing on philosophical ideas about norms, psychological ideas about representations, and most recently neurological ideas about how brains process emotions (Thagard 2006; Thagard and Aubie 2008). The intellectual goal holding all this together is the attempt to build computational models of how the brain produces emotions and uses them in other cognitive processes.

Like research on analogy, it is hard to imagine how theoretical progress on emotion could have proceeded without combining ideas from multiple fields of cognitive science.

16.4 Experimental benefits

Like physics and biology, cognitive science is not a purely theoretical enterprise, but also requires experimental investigations that can be used to evaluate competing theories. Interdisciplinary collaboration has contributed to experimental work in psychology in two ways: through suggesting new kinds of experiments to test interesting theoretical ideas, and through providing new measurement tools for performing experiments.

In the 1960s, the young field of cognitive psychology evolved by developing new kinds of experimental techniques. The growing availability of computers made it much easier to perform experiments that measured the reaction times of subjects performing complex tasks, and the resulting data were used to test the information-processing models of thinking suggested by the new computational theories of mind. The computational models of analogy generated new experimental work to test their predictions. Linguistics also provided new theoretical ideas through Chomsky's work on rules and representations, which inspired new kinds of experiments in psycholinguistics (Pinker 1994). Philosophical ideas have sometimes suggested psychological experiments, as in Rosch's experiments on prototypes. A huge line of experimental research in developmental psychology concerning the ability of children to understand false beliefs originated with philosophical ideas about intention (Boden 2006, p. 488).

In recent years, experiments in cognitive psychology have been most influenced by developments in neuroscience. Ideas about how the brain works have suggested valuable new experiments, but even more importantly neuroscience has provided a whole new set of tools for measuring mental activity. The 1980s saw the development of powerful machines for scanning brains using techniques such as functional magnetic resonance imaging (fMRI) and positron emission tomography (PET). It has become common for cognitive psychologists not only to measure the behavior of experimental subjects when they are performing various tasks, but also to scan their brains while the performance is taking place. Different scanning techniques provide different kinds of detail about the brain regions and temporal courses of neural operations. It is even possible to temporarily disrupt neural processing using transcranial magnetic stimulation. Information about neural processes is also sometimes obtainable by implanting electrodes deep in the brain to stimulate particular regions. Thus, the field of cognitive psychology has been transformed in recent years by the development of new experimental techniques made possible by neuroscience.

Science is most powerful when theoretical ideas mesh with experimental ones; such meshing is very apparent in current attempts to use computational models of brain operations to explain the results of many different kinds of brain scanning experiments. By combining ideas and techniques from psychology, computer science, and neuroscience, cognitive science is successfully pursuing fundamental questions about how the brain works. Answers to these questions are directly relevant to ancient philosophical questions about how minds know reality, make judgments about right and wrong, and appreciate the meaning of life. For example, Thagard (2010) uses psychological and neurological research about vital human needs to argue that the meaning of life is love, work, and play.

Other practical applications include the prospect of improving education by a deeper understanding of the neural mechanisms by which people learn (Posner and Rothbart 2007). The rapidly emerging interdisciplinary field of neuroeconomics is using new knowledge about how brains make decisions to identify the causes of good and bad decisions (Camerer *et al.* 2005). Similarly, political decisions such as voting choice can be illuminated by investigations in psychology and neuroscience (Westen 2007).

16.5 Lessons

The successes and attractive prospects of cognitive science can be attributed to five factors: ideas, methods, people, places, and organizations (Thagard 2005a). It is only useful for people from different disciplines to try to collaborate if there are theoretically powerful ideas that cross disciplinary boundaries. For cognitive science, the main integrative ideas have been representation and computation, which can illuminate the nature of thinking in ways that are useful for all fields of cognitive science—psychology, neuroscience, artificial intelligence, philosophy, linguistics, and anthropology. A representation is a mental structure that can stand for things and events in the world, and inference is a computational mental process that transforms representations. There are other more specific ideas that find valuable applications in many fields, for example particular kinds of representations such as rules and concepts. For instance, some psycholinguists hold that knowledge of language consists primarily of rules such as 'To put an English verb in the past tense, add -ed'. For cognitive scientists, a concept is not a word or an abstract entity, but a mental representation with complex internal structure (Murphy 2002).

In addition, successful interdisciplinary collaboration requires complementary methods. Cognitive science employs many different methods, including psychological experiments, neurological experiments, computer simulations, conceptual analysis, linguistic theorizing, and ethnography. Few people have the time and aptitude to master more than one or two of these methods, but cognitive science benefits from the ways in which methods can be combined to help develop and evaluate explanatory theories about how the mind works. For example, a theory about the nature of concepts can be evaluated on the basis of all of the following: psychological experiments about how people form new concepts; neurological experiments about multiple brain areas involved in the use of concepts; computer simulations of concept learning and application; philosophical reflection on how concepts attach to the world; linguistics studies of concepts in different languages; and ethnographic studies that compare concepts such as color across different cultures. The goal of cognitive science is to arrive at theories that are strongly supported by evidence acquired through all these methods.

The initiation and progress of an interdisciplinary enterprise requires the participation of extraordinary people with the energy and vision to combine the insights of multiple fields. The origins of cognitive science in the 1940s and 1950s benefited from the efforts of exceptional intellectual talents such as Alan Turing, Herbert Simon, George Miller, Noam Chomsky, and Marvin Minsky. Each of these thinkers combined powerful theoretical ability with an appreciation of the insights and methods provided by a variety of different fields. The development of cognitive science organizations in the late 1970s depended on the intellectual vision and organizational skills of another generation of interdisciplinary talents, including Allan Collins, Donald Norman, and Roger Schank. Today, cognitive science depends on a host of people who are active both intellectually and practically in organizations such as the Cognitive Science Society.

Ideas, methods, and people cannot operate in isolation from each other, and occasional conferences are not sufficient to bring about the theoretical and experimental benefits possible from interdisciplinary research. It is therefore important to have places where disciplines can come together on a much more regular basis, at universities or other research institutions. In the 1960s, the Center for Cognitive Studies at Harvard, led by George Miller and Jerome Bruner, brought together many of the early contributors to the interdisciplinary study of mind. Carnegie Mellon University also provided a lively center of activity because of the presence of Herbert Simon and Allen Newell. In the 1970s, other universities such as Yale, Pennsylvania, Berkeley, Michigan, and Edinburgh developed active cognitive science programs, and by the beginning of the twenty-first century there were many places that played the crucial role of fostering such interdisciplinary work. Some do so by explicitly having cognitive science programs, but there are many other related enterprises with different names, such as Harvard's Mind/Brain/Behavior initiative.

Finally, the successful pursuit of an interdisciplinary field is greatly helped by the development of organizations that foster the communication of ideas and methods across fields. For cognitive science, the main organization is the Cognitive Science Society, which began in 1979 and is now complemented by smaller societies operating more locally in Europe and Asia. There also are more specific organizations operating at the intersection of particular pairs of fields, such the Society for Philosophy and Psychology, the Cognitive Neuroscience Society, and the International Conference on Cognitive Modeling. The Cognitive Science Society holds annual conferences that bring together people from many institutions and fields, although psychologists are by far the most heavily represented. The Cognitive Science Society publishes the journal *Cognitive Science* and the new *Topics in Cognitive Science*, which are complemented by a host of other interdisciplinary journals as well as a huge range of periodicals in the various fields of cognitive science. Thus organizations such as societies and journals are an important part of the flourishing of an interdisciplinary field. Goldstone and Leydesdorff (2006) used citation patterns to show that *Cognitive Science* plays a unique bridging role in transferring information across psychology, computer science, neuroscience, and education. Interdisciplinarity can be measured not only by the number of articles produced by multidisciplinary teams, but also by the role that publications play in connecting fields, thereby merging perspectives, tools, and methods.

Like narrower fields, interdisciplinary ventures are far from static, but benefit from changes in ideas, methods, people, places, and organizations. Much cognitive science work has shifted dramatically in recent years toward neuroscience, as many researchers see the study of the brain as providing much of the currently most exciting work on cognition. But not all psychologists, philosophers, or other practitioners share this view, which is just as well. The last thing needed by an interdisciplinary field, or any particular discipline for that matter, is a monolithic approach that narrows down to only a small set of ideas or methods.

In contrast, the full benefits of interdisciplinarity require integration, interaction, and blending of ideas and methods, not their mere juxtaposition and sequencing as found in multidisciplinarity (Klein, Chapter 2 this volume). Cognitive science is sufficiently mature to have its own textbooks, but some are still structured sequentially, describing separately the approaches taken by philosophy, psychology, neuroscience, linguistics, and artificial intelligence (Stillings *et al.* 1995; Friedenberg and Silverman 2006). In contrast, Thagard (2005b) discusses issues about mental representation in an integrated manner that intertwines issues and contributions from different disciplines.

16.6 Conclusion

This chapter has tried to show succinctly how the intellectual benefits of interdisciplinary research can dramatically outweigh the personal and social difficulties of operating in more than one field. Cognitive science provides an excellent illustration of the theoretical and experimental advantages of leaping beyond the confines of particular disciplines. The project of trying to understand the nature of mind is inherently interdisciplinary, requiring the ideas and methods of many different fields. There is still a place for researchers who prefer to restrict themselves to a narrow set of intellectual tools, but progress, especially of the most dramatic sort, requires the mingling of concepts, hypotheses, and methodologies from multiple disciplines. The human brain is so astonishingly complex that we should expect not decades but centuries of ongoing investigations in cognitive science.

Acknowledgements

I am grateful to Robert Frodeman and an anonymous referee for comments on an earlier draft, and to the Natural Sciences and Engineering Research Council of Canada for funding.

References

Appiah, A. (2008). *Experiments in ethics*. Cambridge, MA: Harvard University Press.

Bechtel, W. (2008). *Mental mechanisms: philosophical perspectives on cognitive neuroscience.* London: Taylor and Francis.

Bobrow, D.G. and Collins, A. (eds) (1975). *Representation and understanding: studies in cognitive science*. New York: Academic Press.

Boden, M.A. (2006). *Mind as machine: a history of cognitive science*. Oxford: Clarendon.

Camerer, C., Loewenstein, G.F., and Prelec, D. (2005). Neuroeconomics: how neuroscience can inform economics. *Journal of Economic Literature* **34**, 9–64.

Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.

Dayan, P. and Abbott, L.F. (2001). *Theoretical neuroscience: computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.

Dreyfus, H. (1979). *What computers can't do*, 2nd edn. New York: Harper.

Friedenberg, J. and Silverman, G. (2006). *Cognitive science: an introduction to the study of mind*. Thousand Oaks, CA: Sage.

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

Gentner, D. (1983). Structure mapping: a theoretical framework for analogy. *Cognitive Science* **7**, 155–70.

Gentner, D., Holyoak, K.H., and Kokinov, B.K. (eds) (2001). *The analogical mind: perspectives from cognitive science*. Cambridge, MA: MIT Press.

Goldstone, R.L., and Leydesdorff, L. (2006). The import and export of cognitive science. *Cognitive Science* **30**, 983–93.

Hebb, D.O. (1949). *The organization of behavior: a neuropsychological theory*. New York: Wiley.

Hesse, M. (1966). *Models and analogies in science*. Notre Dame, IN: Notre Dame University Press.

Holyoak, K.J. and Thagard, P. (1995). *Mental leaps: analogy in creative thought*. Cambridge, MA: MIT Press/Bradford Books.

Knobe, J. and Nichols, S. (2008). *Experimental philosophy*. Oxford: Oxford University Press.

Kuhn, T.S. (1970). *The structure of scientific revolutions*, 2nd edn. Chicago: University of Chicago Press.

McCorduck, P. (1979). *Machines who think*. San Francisco: W. H. Freeman.

McCulloch, W.S. (1965). *Embodiments of mind*. Cambridge, MA: MIT Press.

Miller, G.A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review* **63**, 81–97.

Murphy, G.L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Pinker, S. (1994). *The language instinct: how the mind creates language*. New York: William Morrow.

Posner, M.I. and Rothbart, M.K. (2007). *Educating the human brain*. Washington, DC: American Psychological Association.

Rosch, E.B. and Mervis, C.B. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology* **7**, 573–605.

Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences* **3**, 417–24.

Simon, H.A. (1991). *Models of my life*. New York: Basic Books.

Stainton, R. (ed.) (2006). *Contemporary debates in cognitive science*. Malden, MA: Blackwell.

Stillings, N.A., Weisler, S.E., Chase, C.H., Feinstein, M.H., Garfield, J.L., and Rissland, E. (1995). *Cognitive science: an introduction*, 2nd edn. Cambridge, MA: MIT Press.

Thagard, P. (1992). *Conceptual revolutions*. Princeton, NJ: Princeton University Press.

Thagard, P. (2005a). Being interdisciplinary: trading zones in cognitive science. In: S.J. Derry, C.D. Schunn, and A. Gernsbacher (eds) *Interdisciplinary collaboration: an emerging cognitive science*, pp. 317–39. Mahwah, NJ: Erlbaum.

Thagard, P. (2005b). *Mind: introduction to cognitive science*, 2nd ed. Cambridge, MA: MIT Press.

Thagard, P. (2006). *Hot thought: mechanisms and applications of emotional cognition*. Cambridge, MA: MIT Press.

Thagard, P. (ed.). (2007). *Philosophy of psychology and cognitive science*. Amsterdam: Elsevier.

Thagard, P. (2010). *The brain and the meaning of life*. Princeton, NJ: Princeton University Press.

Thagard, P. and Aubie, B. (2008). Emotional consciousness: a neural model of how cognitive appraisal and somatic perception interact to produce qualitative experience. *Consciousness and Cognition* **17**, 811–34.

Turing, A.M. (1950). Computing machinery and intelligence. *Mind* **59**, 433–60.

Westen, D. (2007). *The political brain: the role of emotion in deciding the fate of the nation*. New York: Public Affairs.

Wiener, N. (1961). *Cybernetics: or control and communication in the animal and the machine,* 2nd ed. Cambridge, MA: MIT Press.

Williamson, T. (2007). *The philosophy of philosophy*. Malden, MA: Blackwell.