THE HISTORY OF COGNITIVE SCIENCE: SEVEN KEY DATES*

Margaret A. Boden (maggieb@cogs.susx.ac.uk)

Abstract: Cognitive science has seen seven key dates. The first four were 1943, 1956, 1958, and 1960. Important things happened later, too: in 1969, 1986, and 1987. Those seven years (with 1947 and 1979 as runners-up) all saw seminal publications and/or influential interdisciplinary meetings, in which different methodologies and research opportunities were introduced or highlighted--or, in one case, trenchantly attacked. The current profile of the field has been shaped accordingly.

Keywords: cognitive science, computational psychology, cybernetics, artificial intelligence, perceptrons, connectionism, artificial life, functionalism

* This paper was the opening keynote speech at a meeting of the European Cognitive Science Society in Delphi (June 2007), and appeared in the conference Proceedings.

I: Introduction

Cognitive science has been studied for some seventy years, and covers six different disciplines: AI/A-Life, psychology, neuroscience, linguistics, anthropology, and philosophy. So distilling its history into just seven dates is highly artificial. (Although it is not wholly arbitrary: I intended to choose only five dates, five being a ’nice’ number, but found myself forced to add another two). The thumbnail sketch that follows is based on the book I have recently written on the topic, where everything mentioned here is explored in much greater detail (Boden 2006).

The thematic heart of cognitive science is psychology, and its intellectual heart is AI/A-Life. In other words, it is the study of *mind as machine,* its core assumption being that the same scientific concepts apply to minds and mindlike artefacts.

Since the machines in question are of two main types, there are two major theoretical pathways across the field. One is grounded in logical-symbolic computation, or GOFAI (Good Old-Fashioned AI). The other features adaptive, self-organizing, and/or feedback devices--including certain sorts of connectionist system. We may call them the cybernetic/connectionist and the symbolic--but they both arose out of the cybernetics movement of the 1940s, as we shall see. The field’s history has been shaped by the contrasts and competition between these two approaches.

II: Wartime Thoughts

The first key date within cognitive science is 1943. That wartime year saw three influential publications. The most important was Warren McCulloch and Walter Pitts’ essay ’A Logical Calculus of the Ideas Immanent in Nervous Activity’.

This combined three hugely exciting, but *prima facie* highly diverse, ideas of the early twentieth century: neurone theory, the Turing machine, and the Russell-Whitehead propositional calculus. The authors argued that these were formally equivalent. That is: every expression of the propositional calculus could be computed by some Turing machine, which in turn could be physically implemented in some definable neural net. Logic, computation, and the brain were all of a piece.

In seeing the mind/brain as a Turing machine, McCulloch and Pitts were not thinking only of cognition: for *all* psychological processes, they said, "the fundamental relations are those of two-valued logic". Even in psychiatry, they added, "Mind no longer goes 'more ghostly than a ghost'". Formal networks should be the psychologist's goal: "specification of the net would contribute all that could be achieved in [psychology, however defined]".

These ideas inspired John von Neumann immediately, leading him to design his computer as a machine grounded in binary (true/false) logic instead of decimal arithmetic. But their influence on theoretical psychology was delayed, for three reasons. First, the paper appeared in an obscure journal which few psychologists saw. Second, it used a rebarbative logical formalism (borrowed from Rudolf Carnap), guaranteed to repel most readers. And last, it had no connection with the various wartime problems dominating psychologists' minds in 1943. Its significance would be widely realized only later.

McCulloch and Pitts here initiated *both* theoretical pathways of cognitive science. On the one hand, their paper led to the psychologically oriented connectionism of the 1940s/1950s, initially implemented in wire-and-solder contraptions, not in general purpose computers. On the other hand, once digital computers arrived a few years later, their paper was seen to imply that language-based meanings and reasoning could be modelled by them. (McCulloch, in fact, had long been a follower of the logical atomists' philosophy of language.) That is, it seemed reasonable to hope both that symbolic AI was possible and that it could be seen as theoretical psychology.

Four years later, these two authors would admit that their precisely structured logical networks, and unvarying neural thresholds, did not reflect the noisy, error-prone, and damageable nature of the brain (Pitts and McCulloch 1947). So they now outlined a statistical form of connectionism--and even suggested which parts of the brain perform which types of computation. As the pioneering paper in computational neuroscience, and in distributed computing and probabilistic networks too, this might tempt one to add an *eighth* key date to the list. However, their later paper did not attract many followers. Moreover, they saw it as an "extension" of the earlier one, whose core claim--that neural networks can be *theoretically* mapped onto binary logic--was specifically repeated. Let's mark 1947, then, as an honourable runner-up.

(A brief digression: One might argue that that year should be included in our list for another reason, namely that it saw the first AI manifesto. This was Alan Turing's (1947) report to the National Physical Laboratory, in which he sketched ideas about symbolic AI, search, heuristics, genetic algorithms, and connectionism too. However, this potentially explosive document was circulated only among the *cognoscenti* in the UK. It was not officially published until 1969, when it appeared in one of the *Machine Intelligence* volumes.)

While McCulloch and Pitts had been writing their ground-breaking paper, three other members of the cybernetic community--including Norbert Wiener himself--had been analysing "purpose and teleology" in terms of negative feedback (Rosenblueth et al. 1943). This, they said, could be used so as to reduce the differences between the current state and the goal--an idea that was mentioned in the 'Logical Calculus' paper, too. As they put it, "The signals from the goal are used to restrict outputs which would otherwise go beyond the goal" (p. 19). The examples they listed included heat-seeking missiles, and the muscular overshoot seen (in grasping a glass, for instance) in Parkinsonism. In general, adaptive 'goal-seeking' behaviour of humans and animals was assumed to be controlled in this way.

However, these authors thought of the "goal" as a target, rather than a goal. (In their most persuasive example, heat-seeking missiles, it was exactly that.) The key was perception, not intention. Goals (and sub-goals) considered as imaginary, and intended, future states were not in question. Nor could they be. For there was no mention of internal models, or representations, of the goal--or of current states of the world.

In the very same year, those very matters were being highlighted across the Atlantic in Kenneth Craik's little book *The Nature of Explanation* (1943). This introduced the notion of cerebral models, borrowed from the neurologist Henry Head, into cognitive psychology and the philosophy of mind (Craik described his book as a work of *philosophy).* And it glossed them, for the first time, in terms of the functioning of man-made machines.

The machines Craik had in mind were analogue devices, such as the tidal predictor and the differential analyser. The representational power of cerebral models, he said, lay in the fact that--like the machines just mentioned--each one was "a physical working model which works in the same way as the process it parallels, in the aspects under consideration at any moment" (1943: 51). And he offered some specific hypotheses about the neurophysiology of various analogue "models" for perception.

Although Craik called his approach "a symbolic theory of thought" and referred to "symbolism" in the brain, he seemed to be thinking of representation in general (including language) rather than the logical-computational variety. He died (in an accident) in 1945, so never saw the rise of GOFAI. Probably, he would have accepted formal-symbolic representations as alternative types of cerebral model. Certainly, many of his followers did. Two early cognitive scientists who acknowledged Craik's inspiration were Richard Gregory (e.g. 1966) and Jerome Bruner (who had visited Craik's group in England in 1955-56).

Largely as a result of these three publications of 1943, the next quarter-century saw pioneers working on both types of AI and/or computational psychology.

It would be misleading to say that they were working on both sides of the theoretical *divide,* because the unpleasantly antagonistic schism between connectionism and GOFAI, or (more broadly) between bottom-up and top-down approaches, had not yet developed. At that time, there was still one intellectual community ("cybernetics"), with shared aims and interests. To be sure, some people were focussing more on adaptation and self-organization, others more on logic and meaning--although a few, such as McCulloch, tackled both. Indeed, the rapid rise of GOFAI was mainly due to its promise, not matched by the adaptationists, to deal with inference and linguistic meaning. In general, however, the two sides communicated freely and agreed to differ on what might be the most promising theoretical approach. Only much later did the community separate into distinct sociological camps, with little love lost between them (see Section V).

III: The 1950s

The key dates of the following decade were 1956 and 1958. Indeed, 1956 was the *annus mirabilis* of cognitive science. It saw no fewer than six events that raised the spirits of the nascent cognitive scientists, convincing them that something exciting was happening. Four were publications, and two were meetings aimed at consciousness raising in the emerging interdisciplinary community.

The publications included a book reporting an imaginative series of psychological experiments: Bruner's *A Study of Thinking* (Bruner et al. 1956). (The title alone was a provocation, in those behaviourist days.) Bruner posited several information-processing strategies for concept learning, each more or less appropriate depending on the circumstances--and all defined in broadly computational terms. His ideas would be reflected in much early AI and computational psychology.

In addition, there were three papers. The most influential was George Miller's (1956) information-theoretic 'The Magical Number Seven'--which by the mid-1970s had become the most-cited paper in the whole of cognitive psychology. Another described the first computer model of Donald Hebb's "connectionism" (the word was his coinage). This showed that Hebbian theory could be implemented, but only if his 'ft/wt' learning rule was expressed more precisely (Rochester et al. 1956). And--across the ocean--the last was Ullin Place's 'Is Consciousness a Brain Process?' (Place 1956).

Place's paper is the outlier here, for his mind-brain identity theory was not a contribution to cognitive science as such: it said nothing about mind as machine. But it was eagerly welcomed by scientifically-minded readers, and its materialist spirit--though not its reductionist letter--was retained when philosophical functionalism replaced it four years later (see Section IV).

As for the two consciousness-raising meetings of 1956, the first was the "Summer Research Project on Artificial Intelligence", at Dartmouth College. Organized by the youngsters Marvin Minsky and John McCarthy, this introduced AI to a wider audience. (It also launched the discipline's name, which has been a philosophical millstone around its neck ever since.)

For instance, Minsky handed out the draft of an insightful review of early AI (Minsky 1956). Published a few years later as 'Steps Toward Artificial Intelligence', this was widely seen as AI's manifesto. (As noted above, Turing's earlier manifesto had not been widely circulated outside the UK.) Or perhaps one should rather say "as GOFAI's manifesto", for it argued that connectionist AI had fundamental limitations not shared by symbolic AI. It did, however, suggest that a *combination* of neural networks and GOFAI would be needed to emulate human thought--a suggestion that went largely unheeded. (Minsky himself seemed to forget it in the 1960s, as we shall see, but he eventually followed it up in his "society" theory of mind.)

The Summer Project was not a meeting in the usual sense, but a two-month period during which about a dozen AI pioneers were located at Dartmouth, and anyone who was interested could drop in. The core group included Arthur Samuel and Oliver Selfridge--and, for the final week, Allen Newell and Herbert Simon. In the earlier weeks they had played truant, trying to finish programming their Logic Theorist. This proved theorems (in propositional logic) from the Russell-Whitehead *Principia Mathematica,* and even found a more elegant proof for one of them (Newell et al. 1957).

The Logic Theorist was not the first AI program, though it is often described that way. Quite apart from 'toy' programs written by Turing and others, Samuel had implemented a heuristic program for playing checkers (draughts) in 1949, and a learning version was up-and-running early in 1955 (Samuel 1959: 72). It had even featured on American TV in February 1956, six months before the Dartmouth event. Unlike Newell and Simon, however, Samuel attended that meeting without bringing along printout evidence. That's partly why the participants were more enthused by the Logic Theorist. In addition, logic struck most people as more impressive--more 'human'--than draughts.

But the main reason why more interest was aroused by Newell and Simon's program was that it was explicitly intended as a model of human thinking, guided by Gestalt psychology and by their own experiments. In their view, computers and psychology should be seen as equal partners: "artificial intelligence was to borrow ideas from psychology and psychology from artificial intelligence" (Newell and Simon 1972: 883). Buffs on both sides of this disciplinary fence were excited accordingly.

The second 1956 meeting was the IEEE's three-day Symposium on Information Theory, convened at MIT in mid-September--almost back-to-back with the Dartmouth event. This had more direct influence in bringing psychologists into cognitive science. For among the papers given there were Miller's 'Magical Number Seven', Newell and Simon on the Logic Theorist, and Noam Chomsky on formal grammars--which showed that language, considered as structured sentences not just as word strings, can be formally described.

Miller himself instantly put those other two talks together: "I went away from the Symposium with a strong conviction, more intuitive than rational, that human experimental psychology, theoretical linguistics, and computer simulation of cognitive processes were all pieces of a larger whole, and that the future would see progressive elaboration and coordination of their shared concerns" (quoted in Gardner 1985: 29). This epiphany soon led him to play a crucial role in establishing cognitive science as such (see Section IV).

So 1956 was a good year for the field--but it was soon followed by another. In November 1958, a four-day interdisciplinary seminar took place at the National Physical Laboratory (NPL) in London--a resonant venue, given its post-war connection with Turing. Hosted by the psychophysiologist Albert Uttley, this brought other leading neurophysiologists--Horace Barlow, for example--into the discussion.

About two dozen people, almost all now important names in cognitive science, gathered there. Most had experience of interdisciplinary thinking, having done war-work on the design and use of various novel machines. And Craik was a highly respected name--and, for several attendees, an inspiring personal memory. Recognized intellectual leaders such as McCulloch, Frederic Bartlett, and the anatomist J. Z. Young were joined by youngsters who today are at least as famous. And the youngsters served up some very rich fare.

The atmosphere was electric: it was clear that something exciting was happening. The importance of this meeting for both "sides" of cognitive science--and for AI, A-Life, psychology, and neuroscience--can be indicated by listing a few of the talks (see Blake and Uttley 1959). Among NPL's many memorable moments were these: Selfridge on *Pandemonium;* Frank Rosenblatt on perceptrons; Barlow on his 'coding' theory of perception; Gregory on the misuse of brain-ablation studies; Donald MacKay on the need for hybrid (analogue-digital) machines; McCarthy on giving programs "common sense" *via* predicate calculus (and Yehoshua Bar-Hillel's critical reply on what's now called the frame problem); Gordon Pask on his electrochemical model of a developing concept; and, not least, Minsky on heuristic programming--who summarized the AI manifesto circulated at Dartmouth two years earlier.

The NPL meeting was only one of three events which made 1958 special. The others were two highly contrasting papers, both published in the same volume of *Psychological Review* and both--at least for a while--hugely influential.

The first to appear was a theory of human problem solving, based on the Logic Theorist and its successor the General Problem Solver, or GPS (Newell et al. 1958). Even more powerful than the Logic Theorist, GPS whetted the appetite of psychologists who had not yet heard of the Logic Theorist, and enthused those who had still further. They were attracted, too, by the programme of ongoing psychological experimentation initiated by the authors.

The second seminal paper was Rosenblatt's (1958) account of "perceptrons", also featured at NPL but here reaching a much wider audience. This described a class of connectionist computer models based on Hebbian theory, and focused not on problem solving but on pattern recognition. They could learn to distinguish an A from a B, for example.

Although perceptrons excited many people, including youngsters entering AI, they did not convince everyone whom one might have expected to be sympathetic. Indeed, when cognitive science's manifesto appeared two years later (see Section IV), they were near-invisible: even in those hope-filled pages, parallel processing would be mentioned only in two footnotes. Rosenblatt's hopes were more robust. He saw perceptrons as prefiguring a general theory of human psychology, and was even more daring--some would soon say even more preposterous--than Newell and Simon in his predictions concerning future versions of his machine.

It's noteworthy that these two papers were published in the same Journal. That might have happened ten years later--but not ten years after that. For by then the schism mentioned in Section II had emerged: the field's two pathways had diverged not only theoretically but sociologically too.

IV: Meeting-House, Manifesto, and Mind

Most of the influences mentioned so far were drawn together in two ground-breaking projects of 1960. One was cognitive science's first research institute, the other its manifesto.

Harvard's provocatively named Center for Cognitive Studies was co-founded by Bruner and Miller. Bruner had been running a seminar on these matters for some years, attended (for instance) by the young Chomsky and Jerry Fodor. That had sown important seeds in the local community, for Chomsky later acknowledged Bruner's (neo-Craikian) influence on his positing inner representations of syntactic structure. But in 1960 the new Center put interdisciplinary cognitive science publicly on the academic map.

The name was provocative because it was rejecting behaviourism, then dominant in US psychology. But the word "Cognitive" carried less weight than is often thought, being used simply as an anti-behaviourist shorthand. As Miller later put it: "[We] were setting ourselves off from behaviorism. We wanted something that was *mental* --but "mental psychology" seemed terribly redundant" (Miller 1986: 210). In speaking of "cognition", he said, they were not intentionally excluding "volition" or "conation", but "just reaching back for common sense". In short, even though in practice most cognitive scientists have focused on cognition, the field has always been concerned in principle with *all* aspects of the mind--as McCulloch and Pitts had urged in 1943.

Besides co-founding the Center, Miller offered another spur to cognitive science in 1960. This was his remarkable book *Plans and the Structure of Behavior,* written with Eugene Galanter and Karl Pribram (MGP for short). The book was (unavoidably) simplistic, and careless to boot. Nevertheless, it was a work of vision. Its declared goal was to discover "whether the cybernetic ideas have any relevance for psychology" (p. 3), and its answer was a confident "Yes".

MGP used the notion of a Plan--simply defined as a TOTE unit (Test-Operate-Test-Exit), or as TOTEs made up of lower-level TOTEs--to sketch mental processes. Their discussion ranged over the whole of psychology. Animals and humans; instinct and learning; language and memory; habit and motor skill; chess and choice; values and facts; self image and social role; knowledge and affect; intention and desire; hope and morality; personality and hypnosis; normal life and psychopathology ... *everything* was included.

Plans was the first book to apply computational ideas so widely. Thanks to the recent work of Newell and Simon and of Chomsky, all of whom were repeatedly cited, the most persuasive parts of the book concerned cognition. But the promises reached beyond the persuasion. Miller and Bruner's intention that the "Cognitive" in "Cognitive Studies" should really be read as "mental"--*anything* mental--was reflected in this volume.

Even sympathetic readers were almost deafened by the sound of handwaving. However, they were excited too. For some years, the book would function as a manifesto for the new science of the mind. (A good way of judging progress in cognitive science is to compare today's achievements with the hopes and promises expressed therein.)

Meanwhile, a mile or so away from the new Center, another 1960 landmark had been constructed: Hilary Putnam's functionalism (Putnam 1960). For budding cognitive scientists, this new philosophy offered relief, revelation, and promise. It escaped various dilemmas that had plagued the philosophy of mind--including Place’s identity theory--through the 1950s. More to the point, it saw Turing computation as the causal process at the core of mental life, and the mind as the ’program’ of the brain. By implication, it underwrote the AI-based theoretical psychology that was already emerging.

There were naysayers, of course. Indeed, competing varieties of functionalism would later develop within cognitive science. And there would be plenty of objections from philosophers outside the field. (Putnam himself rejected it, eventually.) Nevertheless, this paper had given sharp philosophical teeth to those who wished to chew the mind in computational terms.

By 1960, then, the field had visibly got off the ground.

V: A Temporary Glitch

The fifth key date, 1969, marks a publication seen by some people as a step backwards rather than forwards. On that view, the damage caused in 1969 was not mended until some twenty years later.

MGP were not the only ones to be under-impressed by perceptrons: Minsky, with Seymour Papert, had a low opinion of them too. He had already expressed doubts in his ’Steps’ paper. But in the 1960s, when Rosenblatt’s ideas were threatening to grab the graduate students, and the funding, he (and others at MIT) felt that sterner measures were called for. The result was an explosive little book called *Perceptrons: An Introduction to Computational Geometry* (Minsky and Papert 1969).

As the sub-title implied, this was a mathematical critique. Minsky and Papert showed that simple parallel processors could not do certain things, such as recognizing connectedness, which one might have expected them to do--and which the then-current GOFAI programs could do. And they predicted that more complex versions would not be much better. Admittedly, in 1959-60 Rosenblatt had proved that perceptrons could learn to do whatever they could be programmed to do. His proof was allowed to be both valid and "seductive" (p. 14), but--Minsky and Papert argued--it had little practical relevance in face of the combinatorial explosion. What the widely-hailed perceptrons could actually do was highly limited. In short, they were fool’s gold.

After this publication, funding for connectionism virtually stopped. In the USA it had started to dry up already, thanks to the circulation of the (even more vitriolic) draft of *Perceptrons* during the early-mid 1960s, and to Minsky’s close friendship with the key funder at DARPA (Joseph Licklider).

Carver Mead later spoke of "the twenty-year famine" in connectionism (Anderson and Rosenfeld 1998: 141). But, rightly, he did not put all of the blame onto Minsky and Papert’s shoulders. Rather, he blamed the early-1960s "overhype" about perceptrons--to which they had been responding.

VI: A Double Renaissance

Both of our last two key dates mark a new visibility, not a new activity. Namely, the public renaissance of connectionism--more precisely, of parallel distributed processing (PDP)--in 1986, and of A-Life a year later. In each case, the new visibility prompted an explosion of further activity that's still expanding.

Connectionism had not stopped dead in its tracks in 1969. Throughout the 1970s, important work was done on associative memories and distributed representation. However, it was seen as maverick, and was largely ignored. A consciousness-raising meeting was held in La Jolla in 1979 (Hinton and Anderson 1981), but it was highly technical. Although it enthused the afficionados, few newcomers were enticed to join the band.

What mended the damage done to connectionism's reputation by Minsky and Papert's attack was the publication in 1986 of the PDP 'bible' (Rumelhart and McClelland 1986; McClelland and Rumelhart 1986). This was deliberately written, priced, and targetted to attract graduate students away from GOFAI and into the PDP stable. So it did--and it attracted many philosophers too. They valued it because it offered a more plausible account of concepts and conceptual similarity.

Crucially, the bible (alongside some lectures by Stephen Grossberg) also attracted the funding authorities. DARPA organized an urgent five-month review of their past funding policy, which had near-ignored connectionism for two decades. Although Minsky, one of the first invited speakers, refused to withdraw his 1969 criticisms (see Minsky and Papert 1988), the outcome was that DARPA changed their mind. They initiated "a major new program in neural networks beginning in 1989" (DARPA 1988: xxv), and gave Minsky and Papert a coded rebuke: "Neural network research is not new--it is, rather, newly revived from an obscurity and even disrepute which is now understood to have been undeserved" (DARPA 1988: 23). The twenty-year famine was over.

Here, we should note another runner-up for an eighth key date. To do that we have to backtrack seven years, to a masterpiece that paved the way for the connectionist renaissance: Douglas Hofstadter's *Godel, Escher, Bach* (1979). This was an intoxicating document. It wove music, logic, biology, and Alice in Wonderland into a song of praise for AI/A-Life in general, and parallel distributed processing in particular. It became a cult book, winning the Pulitzer prize and appearing in many languages. (It's still much admired: in 1999 the *New Scientist* invited a dozen people to choose a science book from the last quarter-century to take to a desert island, and three chose this one.)

So why not add 1979 to our list without further ado? Well, for all its brilliance, *GEB* did not outline a research programme that others could take up. However, it did raise the profile, and indicate the breadth, of cognitive science for the general public. Without its insightful flamboyance to ease the way, acceptance of the much dryer PDP bible would have been less immediate--and much less wide.

The last key date marks a further intellectual renaissance. In 1987 Christopher Langton organized the first conference on "artificial life", at Los Alamos. A-Life, he said, concerned "life as it could be", not just "life as we know it": abstract, preferably formal, descriptions of life were the goal. More generally, the focus was on self-organization and bottom-up processing, in various domains.

He circulated the invitation widely. In the event, a wide spectrum turned up: biologists, biochemists, physicists, mathematicians, AI researchers, neuroscientists, and philosophers (and the journalists turned up too). They discovered--as Langton had hoped--that, despite the superficial differences, they had been working on closely related issues.

The interdisciplinarity and excitement rivalled the NPL meeting of 1958--and the 1950s Macy meetings of the cybernetics community, too. Indeed, that community was much in people's minds. Ross Ashby, Grey Walter, and Pask were honoured by their A-Life descendants after being near-forgotten for a generation. Now, they are familiar names in cognitive science.

VII: Conclusion

And that, for a while, was that. It's not that nothing went on: cognitive science has continued to advance since 1987. And, increasingly, neuroscientific detail has been brought into formerly a-biological zones. But nothing of comparable historical importance has occurred in the last twenty years.

Or rather, nothing that can be recognized *today* as having equal weight. There's plenty of new work out there that's promising, of course--including some which is truly fascinating, not run-of-the-mill (see Boden 2006: ch. 17). A few of these examples may turn out to be historical high points. As yet, however, it is too early to tell.

References:

Blake, D. V., and Uttley, A. M. (eds.) (1959), *The Mechanization of Thought Processes,* 2 vols. Proceedings of a Symposium held at NPL on 24-27 November 1958. (London: Her Majesty's Stationery Office).

Boden, M. A. (2006), *Mind as Machine: A History of Cognitive Science,* 2 vols. (Oxford: Oxford University Press).

Bruner, J. S., Goodnow, J., and Austin, G. (1956), *A Study of Thinking* (New York: Wiley).

Craik, K. J. W. (1943), *The Nature of Explanation* (Cambridge: Cambridge University Press).

DARPA (1988), *DARPA Neural Network Study: October 1987-February 1988* (Fairfax, Virginia: AFCEA International Press).

Gardner, H. (1985), *The Mind's New Science: A History of the Cognitive Revolution* (New York: Harper Collins).

Gregory, R. L. (1966), *Eye and Brain: The Psychology of Seeing* (London: Weidenfeld and Nicolson).

Hinton, G. E., and Anderson, J. A. (1981), *Parallel Models of Associative Memory* (Hillsdale, N.J.: Lawrence Erlbaum).

Hofstadter, D. R. (1979), *Godel, Escher, Bach: An Eternal Golden Braid* (New York: Basic Books).

McClelland, J. L., Rumelhart, D. E., and the PDP Research Group (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition,* Vol. 2, *Psychological and Biological Models* (Cambridge, Mass.: MIT Press).

McCulloch, W. S., and Pitts, W. H. (1943), 'A Logical Calculus of the Ideas Immanent in Nervous Activity', *Bulletin of Mathematical Biophysics,* 5: 115-133.

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.

Miller, G. A. (1986), 'Interview with George A. Miller', in B. J. Baars (ed.), *The Cognitive Revolution in Psychology* (London: Guilford Press), 200-223.

Miller, G. A., Galanter, E., and Pribram, K. H. (1960), *Plans and the Structure of Behavior* (New York: Holt).

Minsky, M. L. (1956), *Heuristic Aspects of the Artificial Intelligence Problem.* Group Report 34-55 (Lexington, Mass.: MIT Lincoln Laboratories, December). Revised as 'Steps Toward Artificial Intelligence', *Proceedings of the Institute of Radio Engineers,* 49 (1961): 8-30.

Minsky, M. L., and Papert, S. A. (1969), *Perceptrons: An Introduction to Computational Geometry* (Cambridge, Mass.: MIT Press).

Minsky, M. L., and Papert, S. A. (1988), 'Prologue: A View From 1988' and 'Epilogue: The New Connectionism', in *Perceptrons: An Introduction to Computational Geometry,* 2nd edn. (Cambridge, Mass.: MIT Press), viii-xv & 247-280.

Newell, A., Shaw, J. C., and Simon, H. A. (1957), 'Empirical Explorations with the Logic Theory Machine', *Proceedings of the Western Joint Computer Conference,* 15: 218-239.

Newell, A., Shaw, J. C., and Simon, H. A. (1958), 'Elements of a Theory of Human Problem-Solving', *Psychological Review,* 65: 151-166.

Newell, A., and Simon, H. A. (1972), *Human Problem Solving* (Englewood Cliffs, N.J.: Prentice-Hall).

O'Reilly, R. C., and Munakata, Y. (2000), *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain* (Cambridge, Mass.: MIT Press).

Pitts, W. H., and McCulloch, W. S. (1947), 'How We Know Universals: The Perception of Auditory and Visual Forms', *Bulletin of Mathematical Biophysics,* 9 (1947), 127-147.

Place, U. T. (1956), 'Is Consciousness a Brain Process?', *British Journal of Psychology,* 47: 44-50.

Putnam, H. (1960), 'Minds and Machines', in S. Hook (ed.), *Dimensions of Mind: A Symposium* (New York: New York University Press), 148-179.

Rochester, N., Holland, J. H., Haibt, L. H., and Duda, W. L. (1956), 'Tests on a Cell Assembly Theory of the Action of the Brain, Using a Large Digital Computer', *Institute of Radio Engineers Transactions on Information Theory,* 2: 80-93.

Rosenblatt, F. (1958), 'The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain', *Psychological Review,* 65: 386-408.

Rosenblatt, F. (1962), *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms* (Washington, DC: Spartan).

Rosenblueth, A., Wiener, N., and Bigelow, J. (1943), 'Behavior, Purpose, and Teleology', *Philosophy of Science,* 10: 18-24.

Rumelhart, D. E., McClelland, J. L., and the PDP Research Group (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition,* Vol.1, *Foundations* (Cambridge, Mass.: MIT Press).

Samuel, A. L. (1959), 'Some Studies in Machine Learning Using the Game of Checkers', *IBM Journal of Research and Development,* 3: 211-229.

Turing, A. M. (1947), 'Intelligent Machinery'. Report prepared for the National Physical Laboratory, UK. First published in B. Meltzer and D. M. Michie (eds.), *Machine Intelligence 5* (Edinburgh: Edinburgh University Press, 1969), 3-23.