

Historical Reflections


How the AI Boom Went Bust

Fallout from an exploding bubble of hype triggered the real AI Winter in the late 1980s.

IN MY LAST two columns (June 2023 and December 2023) I followed the history of artificial intelligence (AI) as an intellectual brand and sub-field of computer science, from its creation in 1955 through to the end of the 1970s. While acknowledging that AI faced high-profile skepticism from the mid-1960s onward, I argued the 1970s were a time of steady growth for the AI research community. Contrary to popular belief, the “first AI winter” of the 1970s never happened. The 1980s, in contrast, saw the rapid inflation of a government-funded AI bubble centered on the expert system approach, the popping of which began the real AI winter: a two-decade slump. I will tell that story here, but first I want to say something about how the maturation of AI played out in textbooks and in the computer science curriculum.

AI in the Curriculum

AI researchers dominated the first 10 years of ACM’s A.M Turing Award, suggesting AI initially occupied the intellectual high ground of computer science. Looking at the computer science curriculum hints at a different story, in which AI moved from a marginal subject in the initial degree programs of 1960s to a core field by the end of the 1980s. The history of computer science education remains understudied, but we can get a fuzzy sense of developments by looking at the evolution of ACM’s recommended curricula.² These recommendations



have a complex relationship to actual practice. Likely they were most closely followed by mid-tier institutions, able to hire across a range of specialties but less likely than Stanford or MIT to have the confidence to build their own unique models around in-house expertise. The first ACM model curriculum, from 1968, described 22 undergraduate courses, including one on “artificial intelligence and heuristic programming.” As an advanced “methodology” elective this was recommended only for masters’ students and for undergraduates pursuing a concentration

in theoretical computer science (one of six sample concentrations).^a The course description suggested a lack of faith in the intellectual maturity of AI: “As this course is essentially descriptive, it might well be taught by surveying various cases of accomplishment in the areas under study.”

A decade later, the Curriculum ‘78 working group recommended an elective covering “basic concepts and techniques,” in AI with knowledge representation, search, and system

^a <https://bit.ly/47b8cmu>

architecture as the main topics.^b It also recommended coverage of LISP, an AI-focused language, in the core course on data structures and algorithms. AI was edging toward the mainstream of a rapidly expanding major. Some 15,121 bachelor's degrees in computer science were awarded in the U.S. in 1980–1981 versus just 2,388 a decade earlier.^c

In 1988, an ACM task force chaired by Peter Denning released a report on the computer science curriculum, which identified artificial intelligence and robotics as one of nine core areas.^d ACM's next detailed model curriculum, released in 1991 in collaboration with the IEEE Computer Society, codified AI and robotics as one of 10 top-level subject areas to be covered by all students (albeit with just nine lecture hours, on a par with databases, human-computer interaction, and numerical computation).^e

The gradual mainstreaming of AI in the computer science curriculum was already apparent in the early 1990s when I studied computer science. The University of Manchester offered specialized AI undergraduate and graduate courses, supported by a team of four AI faculty, several allied faculty focused on formal methods and logic, and a cluster of postdocs and funded Ph.D. students. None of them won ACM A.M. Turing Awards or received gigantic grants, but the group's professor had been a student of Herb Simon, and I had the sense of being competently inducted into a well-established body of techniques. Jumping forward to the present day, the Association for the Advancement of Artificial Intelligence has joined ACM and the IEEE Computer Society as a third partner in the latest computer science curriculum update.

The growth of undergraduate AI courses reflected the new availability of textbooks, replacing teaching anthologies with more coherent volumes that attempted to draw out principles and theories. I identified seven AI textbooks published from 1971 to 1977.^{1,7,8,11,12,16} The books reflected and reinforced the exceptional ability of MIT and Stanford

Early AI had imagined general-purpose reasoning engines driven by collections of individual facts.

to shape the AI brand by determining the topics and approaches to be taught elsewhere. Their eight authors all held degrees from MIT or Stanford; three had earned Ph.D.s under the direction of Marvin Minsky. At the time their books were published, four authors worked at the Stanford Research Institute (which had by then separated from the university). The most widely adopted of the early textbooks was published in 1977 by Patrick Henry Winston, the longtime director of MIT's AI lab.¹⁸ Fifteen years later, as a student, I was assigned an updated edition. Winston's first serious competition came from Nils Nilsson, an SRI researcher and eventual Stanford professor, whose text *Principles of Artificial Intelligence* appeared in 1980. Elaine Rich was a recent Ph.D. graduate of Carnegie Mellon when her textbook appeared in 1983. Through several editions with new coauthors it became the main rival to Winston's book.

The major textbooks of the era dealt entirely with symbolic approaches to AI, neural networks having been purged from the mainstream of computer science. Winston never mentioned connectionist approaches even though his book reflected his specialization in machine learning and computer vision, two areas that have today become synonymous with neural networks. Rich dismissed connectionism in two sentences: "Although there have been many attempts to build learning programs starting with a random network, none of them have met with any degree of success. For this reason, we will not discuss this approach any further here."¹³ The techniques we practiced in Manchester were dominated by symbolic AI and expert systems, though we were told about statistically based techniques for natural language parsing. I

took four AI courses without learning anything about neural networks or genetic algorithms, which were confined to a final-year elective."

From Reasoning to Knowledge

Insider histories of AI agree the crucial intellectual shift of the late 1960s and 1970s was a shift of emphasis away from the hunt for powerful reasoning mechanisms and toward more effective ways of representing knowledge. As Rich wrote in her 1983 textbook, "one of the few hard and fast results to come out of the first 20 years of A.I. research is that *intelligence requires knowledge*."¹³

Early AI theorists had imagined general-purpose reasoning engines driven by collections of individual facts. But researchers concluded that a vast amount of background knowledge was needed to accomplish apparently basic tasks, such as correctly parsing out the verbs and nouns in a sentence or understanding a simple dialogue. From 1974 onward, Minsky talked about the idea of using frames to represent types of objects and events in hierarchies. Frames combined procedures, default values, and facts. The approach strongly paralleled object-oriented programming, developed around the same time. I remember learning about ideas such as inheritance and subclassing in my AI classes rather than my programming courses.

Under Minsky's direction, a generation of researchers trained at MIT worked on microworlds. Searching through a tree of possible states for a desired goal was still the central mechanism in AI, but any general-purpose AI system would confront so many possible sequences of actions that a computer would run out of time and memory long before settling on a reasonable decision. Restricting the complexity of the modeled world made things tractable. The most famous of these systems, and one featured prominently in AI textbooks for decades to come, was SHRDLU, created by Terry Winograd for his 1971 thesis. Winograd's thesis created such a stir that it was published the next year as a full issue of the journal *Cognitive Psychology*.

SHRDLU was described as a program for understanding natural language. It accepted English language questions and commands submitted

b <https://bit.ly/48i87yy>

c <https://bit.ly/3RsTP7v>

d <https://bit.ly/48i8c1Q>

e <https://bit.ly/3GPv40e>

via teletype and typed out responses to the user. The microworld it simulated was a table littered with blocks of different shapes, sizes, and colors that could be placed on top of each other by an imaginary robot arm. The computer's console display rendered the block world in wireframe graphics. The extreme simplicity of the simulated world let Winograd integrate parsing and modeling, implementing each verb as a subroutine. In a lengthy dialogue, SHRDLU responded politely and correctly to questions such as "Is there anything which is bigger than every pyramid but is not as wide as the thing that supports it?" It could answer questions about its own actions, flag ambiguities in questions, and correctly resolve pronouns.

For decades to come, anyone studying AI was likely to learn about SHRDLU and to read an extract from the famous dialogue between Winograd and his creation. But SHRDLU also encapsulated the limitations of traditional AI. While textbook authors looked for unifying principles, most notably search techniques and knowledge representation, the continuing intractability of the key problems addressed by AI researchers meant that textbooks consisted mostly of detailed descriptions of highly specialized systems, few of which were ever applied beyond carefully chosen demonstration problems. SHRDLU's dazzling demonstration script exemplified this, by giving the illusion of having achieved far more than it actually had. As Michael Wooldridge put it, researchers expected "that the techniques it embodied might provide a route to more general natural-language understanding systems, but this hope was not realized."¹⁹ Winograd later became a critic of his own early work, saying the impressive dialogue had been carefully scripted and that even within its limited domain his program was never robust enough to work reliably.¹⁷ He turned away from AI research, becoming instead a theorist of software design and human-computer interaction.

Expert Systems

Although theoretical computer science had displaced AI as the most fertile ground for Turing Awards, the prize committee returned to the field

in 1994 to honor a second generation of AI researchers with awards to Edward A. Feigenbaum and Raj Reddy. Reddy, a pillar of Carnegie Mellon's AI program, had built startlingly capable speech recognition systems, based on a model of separate processes using a blackboard to exchange information.

Feigenbaum's Turing Award profile introduces him as the "father of expert systems," a brand that in the 1980s was often promoted as a less controversial alternative to artificial intelligence.^f Feigenbaum, a Stanford professor and student of Herb Simon, launched the Heuristic Programming Project in the late 1960s. Like Minsky and many other AI researchers, Feigenbaum emphasized the importance of encoding knowledge. But his focus was on automating the work of human experts, initially scientists and doctors. His first system, Dendral, was developed in collaboration with Nobel prize-winning scientist Joshua Lederberg to guess the structure of chemical compounds when fed with formulae and mass spectrogram data.

Feigenbaum and his graduate students went on to develop many other expert systems, including Mycin, a tool for the diagnosis of blood infections. This led in turn to Emycin, which extracted the core reasoning part of Mycin to create a shell that could be loaded with rules encoding expert knowledge from other domains. Distilling expert knowledge into rules was the work of skilled *knowledge engineers*. First they interviewed experts, then they formulated candidate rules. Loading these rules into an inference

f <https://bit.ly/4aFIVEe>

Replacing scarce and expensive human experts with packages of rules was a compelling pitch.

engine such as Emycin and running them against test cases let the knowledge engineer see where it made mistakes, then explore the chain of rules that led to the error and consult the expert to determine what needed to be changed. Soon, claimed Feigenbaum, the system works as well as a human expert. Recent AI approaches involve training systems automatically against huge volumes of data. Feigenbaum insisted (and still insists) that expert systems need only a few hundred carefully chosen rules to equal the decision-making ability of high-functioning professionals.

When the Boom Was On

Replacing scarce and expensive human experts with packages of rules was a compelling pitch. Expert systems launched a wave of private investment in AI, with startup companies selling software tools, system-building services, application-specific services, and implementations of the Lisp and Prolog programming languages. Apparent proof that expert systems could save money in practice was provided by the XCON system designed by Carnegie Mellon professor John McDermott to automate the translation of customer requirements for DEC's VAX computer systems into manufacturing configuration. The initial release condensed expert knowledge into 480 configuration rules, implemented using a specialized language developed with DARPA funds.^g Almost every textbook or magazine discussion of expert systems explained that XCON had eliminated a lengthy review and testing process to shorten VAX delivery times by months. DEC boasted that XCON and a related system saved more than \$40 million a year.

Startup companies proliferated. MIT alone spawned two companies selling expensive workstations with custom processors designed to run Lisp efficiently. The career of Peter Hart, who I mentioned earlier as one of the creators of the A* search algorithm, captures the ups and downs of AI. When ARPA money for SRI's robot project dried up, he made a name for himself in expert systems research with the Prospector geological sys-

g <https://bit.ly/3vaEa1N>


tem, then ran an AI lab for Schlumberger Ltd., and in 1983 partnered with fellow SRI veteran Richard Duda to start an expert system services company named Syntelligence. McDermott too founded a company, the Carnegie Group. Feigenbaum himself cofounded three companies. As Hart recalled the era, “new expert systems companies were being formed at the rate of what seemed like one a week.”⁶

Like the earlier waves of AI enthusiasm, the new boom had a lot to do with government spending. This time it was fear of Japan, rather than the USSR, that unlocked the public purse. Japan’s commitment to a human-centered approach to computing in its high-profile Fifth Generation Project included an effort to create natural language interfaces. Feigenbaum led a hugely successful campaign to present this as a major economic threat to the U.S., warning that only massive public investment in expert systems could prevent Japan overtaking the U.S. in computing just as it had in television and motorcycle manufacturing. Feigenbaum called for “a national plan of action, a kind of space shuttle program for the knowledge systems of the future.”^{4,5}

Politicians attempted to capitalize on a widespread belief that a microcomputer revolution was about to usher in a post-industrial society or information society in which leadership in computer technology would be much more important than traditional manufacturing industry as a contributor to national success. Britain launched the Alvey project and Europe established the transnational ESPRIT research initiative.

The most ambitious project of the era was Cyc, led by former Stanford and Carnegie Mellon faculty member Doug Lenat, a specialist in systems that made discoveries. Whereas expert systems aimed to capture knowledge in extremely narrow domains, Lenat dreamed of equipping an AI logic engine with an everyday knowledge base broad enough that it could add automatically to its base of facts and even invent new heuristics. That would take a lot of knowledge: the Cyc name came from “encyclopedia.” Lenat estimated codifying an encyclopedia worth of knowledge into a gi-

Google’s Ngram Viewer, based on a large English text corpus, suggests discussion of AI surged in the 1980s, driven by interest in expert systems, but declined throughout the two-decade “AI winter” that followed. Source: <https://bit.ly/3ROXigO>



gantic semantic network would take approximately 2,000 years of person-effort. After that, the system would know enough to assimilate everything else by reading books and newspapers. Starting in 1983, Lenat got 400 researchers and more than \$500 million from the Microelectronics and Computer Technology Corporation (MCC), an industrial consortium sponsored by the U.S. government to counter the Japanese threat.

The AI Winter

DARPA jumped back into AI in a big way in 1983 with its Strategic Computing Initiative, the story of which was told in a fascinating book by Alex Roland and Philip Shiman.¹⁴ The program was sold to Congress with promises of direct military applications, and rested on the assumption that existing approaches to expert systems, natural language understanding, and vision were ready for large-scale application once computer hardware improved (something the program aimed to accelerate with support for research on massively parallel supercomputers, microelectronics, and prototyping). These technologies would be integrated into military systems, with self-driving vehicles selected as a test case.

In 1984, a distinguished panel convened at the annual meeting of the American Association for the Advancement of Artificial Intelligence.

The conference was starting to feel like a trade show. Expert system startups were mushrooming, large corporations were rushing to establish AI groups, government money was flooding in, and a frenzied job market ensured lucrative employment for anyone who could claim a few months of AI experience. Yet introducing the panel on “The Dark Ages of AI,” Yale professor Drew McDermott warned of a feeling of “deep unease” that excessively high expectations for AI “will eventually result in disaster.” To sketch a worst-case scenario,” continued McDermott, “suppose that five years from now the strategic computing initiative collapses miserably as autonomous vehicles fail to roll. The Fifth Generation turns out not to go anywhere, and the Japanese government immediately gets out of computing. Every startup company fails. Texas Instruments and Schlumberger and all other companies lose interest. And there’s a big backlash so that you can’t get money for anything connected with AI. Everybody hurriedly changes the names of their research projects to something else.”¹⁰

McDermott noted this “unlikely” scenario was so apocalyptic that it was “called the ‘AI Winter’ by some,” in reference to scientific debate over the prospect that nuclear war would throw enough soot into the atmosphere to trigger devastating global cooling in a nuclear winter. Super-

power diplomacy staved off the nuclear winter but by the end of the decade the AI apocalypse was taking place just as described.

At DARPA, for example, speech recognition work progressed well but other strategic computing projects disappointed. Reagan-era budget cuts also contributed to a scaling back of effort and expectations. At the end of 1987, DARPA abandoned the flagship effort to build an autonomous land vehicle (though work it had funded at Carnegie Mellon's Navlab provided an important foundation for later developments). DARPA's leadership "elected simply to sweep Strategic Computing under the carpet and redirect computer research toward the 'grand challenges' of high-performance computing. Numerical processing replaced logical processing as the defining goal."¹⁴

The AI Winter is clearly visible on Google's Ngram chart in this column. Discussion of AI grew steadily through the 1970s before spiking in the 1980s. This was tied to an explosion of discourse about expert systems, a phrase that at its peak in the late 1980s was just as common as "artificial intelligence" itself. Both fell precipitously during the 1990s. By 2010, references to AI were coming less than one-third as often as they had at the peak and the rate was still falling.

Discussion of expert systems dropped more rapidly, reflecting the collapse of the short-lived industry. Comparing the expert system story with the approximately contemporaneous commercialization of relational database management systems is instructive. Both began with bold ideas of disputed practicality, followed by impressively engineered prototype systems produced in industrial and academic labs. Both technologies were recognized with Turing Awards, and both were commercialized as software platforms marketed by startups with close connections to universities. In the case of relational database management systems, the crucial work was done by IBM Research and at the University of California, Berkeley. Relational database management companies thrived, turning their products into universal infrastructures for corporate data.

The best known of them, Oracle, is among the world's most successful businesses.

In contrast, the market for expert system software proved unsustainable because most companies struggled to build the in-house skills needed to use them effectively. Companies that had set up AI groups and purchased expert system software discovered systems designed to automate expertise required them to hire new experts to maintain them. By 1989, DEC had 59 technical staff members assigned to maintain the infrastructure and base of rules for its internal expert systems, which remained the most widely publicized application of AI.^h Few companies could sustain such investments, particularly as a shortage of AI specialists had driven up wages.

Lenat's grand vision for Cyc did not materialize either, in part because developing a single consistent knowledge base proved impossible, but the project continued. In 1994, as the MCC began to implode, the Cyc project was transferred to a private company that continues to develop and license Cyc. It has now grown to a collection of 30 million rules.^{3,9}

The AI Winter extended to the Turing Awards. In the eyes of 16 successive selection committees, the field of AI failed to produce anything between 1995 and 2010 to match the advances in areas such as databases, cryptography, networking, programming, and complexity theory that were honored with awards.

Broad-based and sustained as this decline in discussion of AI was, it may not reflect experiences outside the U.S. and U.K. and likely understates the resilience of AI as an area of computer science teaching and research. In South Korea, for example, AI publications and funding rose steadily in the late 1980s and early 1990s.¹⁵ Because conventional histories of AI (at least those in English) have constructed AI as an almost entirely Anglo-American project, this and other aspects of its history must be reassessed when that focus eventually broadens.

AI returned to primetime in the 2010s with the dramatic revival of

interest in connectionist approaches centered on deep learning systems. The effort began in the 1980s but, because AI had been redefined around symbolic approaches, was pursued under other brands, such as machine learning and pattern recognition. Only in the last few years has the AI brand itself been flipped to refer primarily to deep learning and generative systems. In my next column I will be telling that story and looking at differences and parallels between our current wave of AI hype and the booms and busts of years gone by. **C**

References

1. Duda, R. and Hart, P. *Pattern Recognition and Scene Analysis*. Wiley, New York, 1973.
2. Dziallas, S. and Fincher, S. The history and purpose of computing curricula (1960s–2000s). In *Communities of Computing: Computer Science and Society in the ACM*, T.J. Misa, Ed. Morgan & Claypool (2017).
3. Eklia, H.R. *Artificial Dreams: The Quest for Non-Biological Intelligence*. Cambridge University Press, New York, 2008.
4. Feigenbaum, E.A. and McCorduck, P. *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*. Reading, MA, 1983.
5. Garvey, C. Artificial intelligence and Japan's Fifth Generation: The information society, neoliberalism, and alternative modernities. *Pacific Historical Review* 88, 4 (2019).
6. Hart, P.E. An artificial intelligence odyssey: From the research lab to the real world. *IEEE Annals of the History of Computing* 44, 1, (Jan.–Mar. 2022).
7. Hunt, E.B. *Artificial Intelligence*. Academic Press, New York, 1975.
8. Jackson, P.C. *Introduction to Artificial Intelligence*. Petroselli Books, New York, 1974.
9. Lenat, D. Creating a 30-million-rule system: MCC and Cycorp. *IEEE Annals of the History of Computing* 44, 1 (Jan.–Mar. 2022).
10. McDermott, D. et al. The dark ages of AI: A panel discussion at AAAI-84. *AI Magazine* 6, 3 (1985).
11. Nilsson, N.J. *Problem-Solving Methods in Artificial Intelligence*. McGraw-Hill, New York, 1971.
12. Raphael, B. *The Thinking Computer: Mind Inside Matter*. W. H. Freeman & Company, San Francisco, CA, 1976.
13. Rich, E. *Artificial Intelligence*. McGraw-Hill, New York, 1983.
14. Roland, A. and Shiman, P. *Strategic Computing: DARPA and the Quest for Machine Intelligence*. MIT Press, Cambridge, MA, 2002.
15. Shin, Y. Hangul and the "spring" of artificial intelligence research in South Korea. *Technology's Stories* 6, 1 (Mar. 2018).
16. Slagel, J.R. *Artificial Intelligence: The Heuristic Programming Approach*. McGraw-Hill, New York, 1971.
17. Winograd, T. *Oral History Interview by Arthur L. Norberg*. Charles Babbage Institute. (1991); <https://hdl.handle.net/11299/107717>
18. Winston, P. *Artificial Intelligence*. Addison-Wesley, Reading, MA, 1977.
19. Wooldridge, M. *A Brief History of Artificial Intelligence: What It Is, Where We Are, and Where We Are Going*. Faltiron Books, New York, 2021.

Thomas Haigh (thomas.haigh@gmail.com) is a professor of history at the University of Wisconsin—Milwaukee, WI, USA, and a Comenius visiting professor at Siegen University, Germany.

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project-ID 262513311-SFB 1187 Media of Cooperation.

^h <https://bit.ly/41v5Wp4>