

Historical Reflections


There Was No ‘First AI Winter’

Despite challenges and failures, the artificial intelligence community grew steadily during the 1970s.

AS I CONCLUDED my June Historical Reflections column, artificial intelligence had matured from an intellectual brand invented to win funding for a summer research workshop to one of the most prestigious fields in the emerging discipline of computer science. Four of the first 10 ACM A.M. Turing Award recipients were AI specialists: Marvin Minsky, Herb Simon, Allen Newell, and John McCarthy. These men founded the three leading AI labs and played central roles in building what are still the top three U.S. computer science programs at MIT, Stanford, and Carnegie Mellon. Conceptually AI was about uncovering and duplicating the processes behind human cognition; practically it was about figuring out how to program tasks that people could do but computers could not. Although *connectionist* approaches based on training networks of simulated neurons had been prominent in the primordial stew of cybernetics and automata research from which AI emerged, all four Turing Award recipients favored the rival symbolic approach, in which computers algorithmically manipulated symbols according to coded rules of logic.

A History of Failed Ideas?

AI was born in hype, and its story is usually told as a series of cycles of fervent enthusiasm followed by bitter disappointment. Michael Wooldridge, himself an eminent AI researcher, began



his recent introduction to the field by remembering when he told a colleague about his plan to tell “the story of AI through failed ideas.” In response, “she looked back at me, her smile now faded. ‘It’s going to be a bloody long book then.’”²²

Major awards lag years behind research. By the time Newell and Simon shared the 1975 ACM A.M. Turing Award the feasibility of their approaches to AI was being increasingly challenged. The AI community would have to wait 19 years for another winner. It was displaced as the intellectual high ground of the emerging discipline by *theoretical computer science*, a field centered on mathematical analysis of algorithms, which garnered nine awardees

during the same period.^a This new focus bolstered the intellectual respectability of computer science with a body of theory that was impeccably mathematical yet unlike numerical analysis, which was falling out of computer science over the same period, not directly useful to or understood by other schol-

^a These awards focused on computational complexity theory and the analysis of algorithms. I am construing theoretical computer science here to encompass the work of Rabin and Scott (1976), Cook (1982), Karp (1985), Hopcroft and Tarjan (1986), Milner (1991), and Hartmanis and Stearns (1993). I am not including winners cited primarily for contributions to programming languages, except for Milner whose citation emphasized theory, though Wirth and Hoare both made important theoretical contributions.

**“I am
human...
just like
you.”**

For the first
time, the AI
had asserted
its claim on
humanity.

At that moment, the AI
and I had become one...

Communications of the ACM is looking for writers in our community to contribute sci-fi short stories, between 1,000 and 1,200 words, for our quarterly “Future Tense” section.

**Do you have
a great story to tell?**
Make contact at
LastByte@cacm.acm.org

Association for
Computing Machinery



ars in established disciplines.

The problems AI researchers had taken as their test cases were difficult in a fundamental mathematical sense that dashed hopes of ingenious breakthroughs. Once AI researchers applied the new techniques of complexity analysis “Everywhere they looked—in problem solving, game playing, planning, learning, reasoning—it seemed that the key problems where NP-complete (or worse).”²² Progress would come slowly and painfully, with methods that worked in some cases but not others.

The early practitioners of AI had consistently and spectacularly overestimated the potential of their methods to replicate generalized human thought. In 1960, for example, Herb Simon had declared “within the near future—much less than 25 years—we shall have the technical capability of substituting machines for any and all human functions in organizations.” He believed the “problem-solving and information handling capabilities of the brain” would be duplicated “within the next decade.” As professionals were replaced by machines “a larger part of the working population will be mahouts and wheelbarrow pushers and a smaller part will be scientists and executives.”²⁰

The same processes of hype that gave AI a high profile for military sponsors and awards committees also made the field a topic of public debate. Promises made for intelligent computers tapped into longer-established myths and science fiction stories of thinking machines and mechanical servants. HAL, the murderous computer from the movie *2001: A Space Odyssey*, whose name was said to be a contraction of *heuristic* and *algorithmic*, was one of many fictional depictions of the promises made by AI researchers. Minsky himself had been hired to get the details right. Meanwhile a series of books appeared criticizing those promises and challenging the feasibility of artificial intelligence.¹⁰

The AI boosters were wrong, of course, though their critics were not always right. Computers had always been sold with futuristic hype, but overly optimistic technical predictions made during the same era for other areas of computer science such as graphics,

computer mediated communication, scientific computation, and databases were eventually met and surpassed. In contrast, the approaches adopted by the AI community conspicuously failed to deliver on Simon’s central promises.

Military Origins of AI

AI began as a Cold War project centered on a handful of well-connected researchers and institutions. The original Dartmouth meeting was funded by the Rockefeller Foundation. A full-scale research program would require deeper pockets, which in the 1950s were usually found attached to military uniforms. When Newell and Simon met and began to collaborate on their famous theorem prover both were employed by the RAND Corporation, a non-profit set up to support the U.S. Air Force. This gave them access not just to RAND’s JONNIAC computer, one of the first modern style computers operational in the U.S., but also to RAND programmer Clifford Shaw who was responsible for squeezing the ambitious program into the machine’s tiny memory.⁵ Frank Rosenblatt developed his perceptrons, the most important of the early neural networks, in a university lab funded by the U.S. Navy. At MIT, Minsky’s early work benefitted from the largess of the Air Force and the Office of Naval Research.

This deep entanglement of early AI with the U.S. military is difficult to overlook. Johnnie Penn highlighted the military dimension in his recent dissertation, challenging the phrase “good old fashioned AI” (a common description of traditional symbolic work) as something that “misrepresents and obscures this legacy as apolitical.”¹⁸ Yarden Katz insists the apparent technical discontinuities in

**This deep
entanglement of
early AI with the U.S.
military is difficult to
overlook.**

the history of AI are just distractions from a consistent history of service to militarism, American imperialism, and capitalism.¹³

Yet AI was not exceptional in this respect. Military agencies supplied approximately 80% of all Federal research and development funding during the 1950s, the first decade of the Cold War. This wave of money flowed disproportionately to MIT and Stanford, which were not only the two leading centers for both AI and computer science but also the primary exemplars of a new, and to many disturbing, model for the relationship between universities, the Federal government, and military needs. Stuart W. Leslie's history book *The Cold War and American Science* focused entirely on those two institutions as prototypes for a new kind of university restructured around military priorities.¹⁴

Computing was, after all, an expensive endeavor and there were few alternative sources of support. Set in the larger picture of military investment in computing, including projects such as the SAGE air defense network and guidance systems for Minuteman missile, the sums spent on AI appear quite small. Most computing projects of the 1940s and 1950s were underwritten directly or indirectly by the U.S. military.⁸ ENIAC, the first programmable electronic computer, was commissioned in 1943 by the U.S. Army for use by its Ballistics Research Laboratory.¹² Such relationships blossomed as the Second World War gave way to the Cold War. IBM, for example, received more money during the 1950s from developing custom military systems such as the hardware for the SAGE air defense network than it did from selling its new lines of standard computer models. And even the market for those standard projects was driven by the Cold War. IBM's first commercial computer model, the 701, was known informally as the "Defense Calculator" and sold almost entirely to government agencies and defense contractors. It was the Federal government, not IBM itself, that managed the delivery schedule for early models to maximize their contribution to national security.¹¹ The needs of military and aerospace projects kick-started the semiconductor industry in what became Silicon Valley.

ARPA supported well-connected research groups without formal competitive review or any commitment to provide specific deliverables.

AI remained heavily dependent on military funding in the 1960s, as labs at MIT and Stanford received generous funding through the newly established Advanced Research Projects Agency. ARPA reportedly spent more on AI than the rest of the world put together, most of which went to MIT and Stanford. Carnegie Mellon was not initially in the same league, but its early success in computing and artificial intelligence won substantial ARPA funding by the 1970s and fueled the rise of the university itself. The National Science Foundation, a civilian agency, was less important. During the 1950s and 1960s it did not have a directorate focused on computer science. It made few grants to support computing research (though it was active in funding computing facilities).³

ARPA supported well-connected research groups without formal competitive review or any commitment to provide specific deliverables. J.C.R. Licklider, the first director of ARPA's Information Processing Techniques Office, joined ARPA from military contractor BBN and had earlier been a member of the MIT faculty. After showering MIT with government money he eventually rejoined its faculty, to run the ARPA-funded Project MAC (into which Minsky and his AI group had been incorporated). Licklider then returned to ARPA for a second term as director. That might all seem a little too cozy by modern standards, but ARPA's early success in fostering emerging computer technologies was spectacular: not just the Internet, which can be traced back to an idea of Licklider's,

but also computer graphics and time-sharing.¹⁷ Paul Edwards summarized the early history of AI in his classic *The Closed World*, arguing that under the influence of ARPA it became "part of the increasingly desperate, impossible tasks of enclosing the U.S. within an impenetrable bubble of high technology." He believed Licklider's vision for interactive computing was shaped fundamentally by military concerns with command and control.⁶

Were the founders of AI who worked at RAND or took money from the Pentagon thereby coopted into an imperialistic effort to project American power globally? Did their work somehow come to embed the culture of the military industrial complex? Historians will likely be arguing these questions for generations to come. AI, like cybernetics, unquestionably benefitted from a powerful alignment with a more general faith of scientific and political elites in abstraction, modeling, and what has been called by historians of science *Cold War rationality*.⁷

Personally, though, I am inclined to see the founders of AI as brilliant boondoggles who diverted a few buckets of money from of a tsunami of cold war spending to advance their quirky obsessions. Steven Levy noted that a "very determined solipsism reigned" among the hackers of Minsky's lab at MIT, even as the antiwar protesters forced them to work behind locked doors and barricades. He quoted Minsky as claiming Defense Department funding was less intellectually corrosive than hypothetical money from the Commerce Department or the Education Department.^{4,16} On the Stanford side, John McCarthy was a proponent of scientific internationalism. He was raised communist and made five visits to the USSR during the 1960s, though his politics drifted rightward in later decades.²¹ Philip Agre, recalling the investments by the military industrial complex in his graduate training at MIT, wrote that "if the field of AI during those decades was a servant of the military then it enjoyed a wildly indulgent master."²

Summers and Winters

When scientists write histories they usually focus on intellectual and technical accomplishments, leaving professional historians and science studies

scholars to raise indelicate questions about the influence of money. In contrast, the insider story of AI as told by experts such as Wooldridge, Nils J. Nilsson and Margaret Boden has been structured explicitly around shifts in government funding.⁴

Why was AI so vulnerable to the whims of government agencies? One factor was the concentration of early AI work in a handful of well-connected labs. Another was the reliance of AI researchers on large and expensive computers. Perhaps the most important was the failure of AI, during its first few decades, to produce technologies with clear commercial potential that might attract a broader range of sponsors. The health of AI as a field thus depended on the ability of researchers to persuade deep-pocketed sponsors that spectacular success was just around the corner.

Relying on a handful of funding sources proved hazardous. Machine translation projects were an early beneficiary of military largess, but this made them vulnerable when their feasibility was questioned. American funding for this area dried up after a damning report was issued in 1966 by the ALPAC committee, a scientific panel sponsored by the Department of Defense, the National Science Foundation, and the CIA to investigate progress in the area.¹⁹

The late 1980s are universally seen

as the beginning of the “AI Winter,” in which faith and funding for AI dwindled dramatically. I will tell that story later, but in a fuzzier way the period from 1974 to 1980 has increasingly been described as an earlier winter for AI.^b This narrative blames the 1973 Lighthill Report, commissioned by the Science Research Council of the U.K., for a collapse of British support for AI work. Across the Atlantic, this is said to have inspired other funders to ask more difficult questions.²²

Sir James Lighthill was commissioned to write his report with the specific intent of justifying the withdrawal of funding for Donald Michie’s lab at Edinburgh, the most important center for AI research in the U.K. Lighthill, an eminent applied mathematician, endorsed both practical work on industrial automation and work to support analysis of brain functions (often called cognitive science) but opposed funding for ambitious work intended to unite the two in a new science of machine intelligence. Jon Agar’s analysis makes clear that the creative blurring of categories which let AI researchers spin nar-

b The narrative of the 1970s as an AI winter seems driven more by online summaries than scholarly history but is becoming firmly entrenched, for example on Wikipedia, which claims that “There were two major winters in 1974–1980 and 1987–1993”; https://en.wikipedia.org/wiki/AI_winter


row achievements into vast promises also left them vulnerable to attacks that challenged implied connections between specific computational techniques and general intelligence.¹

How, and indeed whether, Lighthill’s attack on one controversial lab precipitated a broad international funding collapse for AI remains unclear. It coincided with broader changes in U.S. support for science. In 1973 Congress, inspired by Vietnam-era concerns that the military was gaining too much power over university research, passed legislation explicitly prohibiting ARPA from funding work that was not directly related to military needs. ARPA was renamed to DARPA, the D standing for Defense.¹⁷ As responsibility for basic research funding shifting increasingly to the NSF, DARPA funding required more direct military justification and came with more strings attached.

AI’s Steady Growth in the 1970s

Historical work on AI has so far focused on a handful of elite institutions and researchers, just as I have done in the series so far. Any DARPA-related AI slowdown was felt most deeply at those highly visible sites, whose researchers, and graduates were the people most likely to give keynotes, write memoirs, and shape the collective memory of the discipline. But the big contracts awarded on a handshake only ever went to a few institutions, whereas the institutionalization of AI took place internationally and across a much broader range of universities. In the first major history of AI, Pamela McCorduck noted that 21 years after the Dartmouth conference the influence of its participants and the programs they founded remained strong. Nine invited papers were given at the 1977 International Joint AI Conference: three by Simon and his former students Ed Feigenbaum and Harry Pople; one by Feigenbaum’s own student Doug Lenat; one by Minsky and one by McCarthy. MIT, Stanford, SRI, and Carnegie Mellon dominated, “with the representation from other laboratories being sparser than might have been expected in a field that had grown from the 10 Dartmouth pioneers in 1956 to nearly 1,000 registrants in 1977.”¹⁵

Google Ngram data, based on a large English language corpus and plotted between 1955–1980. References to artificial intelligence rose consistently through the 1970s even as discussion of the related concepts of automata and cybernetics declined sharply.[†]



† For example Wooldridge,²² 61–62.

Computer science developed as a highly federated discipline, in which most practitioners identified and engaged more with their specialist area than with the field as whole. Over time new areas such as networking, databases and graphics gained prominence while others slipped out of the mainstream. AI enthusiasts created SIGART in 1966, one of the first Special Interest Groups (SIGs) within the ACM. The SIGs gave institutional recognition to the subfields of computer science. With their own publications, conferences, and finances they came to account for most of the ACM's activity.

I suspect that judged by metrics such as the number of students enrolled in AI courses, total number of AI researchers, attendance at conferences, or quantity of research publication the story of AI in the 1970s and 1980s would look less like a series of abrupt booms and busts and more like a march toward disciplinary professionalization. As a first step in this analysis, I located two data sources, neither of which supports the idea of a broadly based AI winter during the 1970s.

One is membership of ACM's SIGART, the major venue for sharing news and research abstracts during the 1970s. When the Lighthill report was published in 1973 the fast-growing group had 1,241 members, approximately twice the level in 1969. The next five years are conventionally thought of as the darkest part of the first AI winter. Was the AI community shrinking? No! By mid-1978 SIGART membership had almost tripled, to 3,500. Not only was the group growing faster than ever, it was increasing proportionally faster than ACM as a whole which had begun to plateau (expanding by less than 50% over the entire period from 1969 to 1978). One in every 11 ACM members was in SIGART.

Not all the participants in this growing community worked in elite, DARPA-funded labs. As the *SIGART Bulletin* summed up the AI hierarchy a few years later: "AI research in the U.S. can be viewed as having three major components: a few highly visible sites with major efforts; a large number of sites with smaller numbers of workers; and a diffuse set of researchers and developers in other fields who believe that AI research may be rel-

Not all the participants in this growing community worked in elite, DARPA-funded labs.

evant for them."^c Perhaps the people in those "highly visible sites" were suffering, but general interest in AI continued to grow rapidly.

The other data source is Google's Ngram viewer, which suggests that the term *artificial intelligence* became more common during the so-called AI winter. Its growth stalled for a few years in the mid-1960s but recovered by 1970 and grew steadily through 1980. Yet the mid-1960s are usually described as a golden age for AI.

The AI community founded an independent organization—the American Association for Artificial Intelligence—in 1979. Reporting this news, the SIGART chair Lee Erman noted wistfully "SIGs are set up as arms of ACM and as such must obtain ACM approval for most significant actions, including budgets, new publications, sponsorship of conferences, and interaction with non-ACM organizations. This structure may be appropriate for a 'special interest group' (although I would argue more autonomy would be beneficial to the SIGs and to ACM), but not for a national scientific organization, which needs far more independence."^d This proved prophetic: while SIGART membership continued to grow well into the 1980s, the new association eventually replaced it as the hub of the AI community by developing a panoply of publications, conferences, and awards.

That is all I have space in this column, but in the next installment I will be looking at the codification of AI's intellectual content in early textbooks and its entrenchment in the computer

science curriculum, at the new emphasis in the 1970s on knowledge representation over pure reasoning, and at the spectacular bubble of funding for expert systems in the early 1980s, which burst to create the real "AI Winter." ■

References

1. Agar, J. What is science for? The Lighthill report on artificial intelligence reinterpreted. *British J. for the History of Science* 53, 3 (Sept. 2020), 289–310.
2. Agre, P. Toward a critical technical practice: Lessons learned in trying to reform AI. In *Bridging the Great Divide: Social Science, Technical Systems, and Cooperative Work*. G. Bowker, Ed. Psychology Press, New York, 1997, 131–158.
3. Aspray, W. and Williams, B.O. Arming American scientists: NSF and the provision of scientific computing facilities for universities, 1950–73. *IEEE Annals of the History of Computing* 16, 4 (Winter 1994), 60–74.
4. Boden, M.A. *Mind as Machine: A History of Cognitive Science*. Clarendon Press, Oxford, U.K., 2006.
5. Dick, S. Of models and machines: Implementing bounded rationality. *Isis* 106, 3 (2015), 623–634.
6. Edwards, P.N. *The Closed World: Computers and the Politics of Discourse in Cold War America*. MIT Press, Cambridge, MA, 1996.
7. Erickson, P. et al. *How Reason Almost Lost Its Mind: The Strange Career of Cold War Rationality*. University of Chicago Press, Chicago, IL, 2013.
8. Flamm, K. *Creating the Computer: Government, Industry, and High Technology*. Brookings Institution, Washington, D.C., 1988.
9. Freeman, P.A., Adrien, W.R., and Aspray, W. *Computing and the National Science Foundation, 1950–2016*. Association for Computing Machinery, NY, 2019.
10. Garvey, S.C. The 'general problem solver' does not exist: Mortimer Taube and the art of AI criticism. *IEEE Annals of the History of Computing* 43, no. 1 (2021).
11. Haigh, T. Computing the American way: Contextualizing the early U.S. computer industry. *IEEE Annals of the History of Computing* 23, 2 (Apr.–June 2010), 8–20.
12. Haigh, T., Priestley, M. and Rope, C. *ENIAC In Action: Making and Remaking the Modern Computer*. MIT Press, Cambridge, MA, 2016.
13. Katz, Y. *Artificial Whiteness: Politics and Ideology in Artificial Intelligence*. Columbia University Press, New York, 2020.
14. Leslie, S. *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*. Columbia University Press, NY, 1993.
15. McCorduck, P. *Machines Who Think*. A.K. Peters, Natick, MA, 2004, 130–131.
16. Norberg, A.L., O'Neill, J.E., and Freedman, K.J. *Transforming Computer Technology: Information Processing for the Pentagon, 1962–1986*. Johns Hopkins University Press, 1996.
17. Nillson, N.J. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge University Press, New York, 2010.
18. Penn, N. *Inventing Intelligence: On the History of Complex Information Processing and Artificial Intelligence in the United States in the Mid-Twentieth Century*. Ph.D. dissertation. University of Cambridge, 2020.
19. Poibeau, T. *Machine Translation*. MIT Press, Cambridge, MA, 2017.
20. Simon, H. The corporation: Will it be managed by machines? In *Management and Corporations*. M. Anshen and G.L. Bach, Eds. The McGraw-Hill Book Company, New York, 1960, 17–55.
21. Tatarchenko, K. Transnational mediation and discipline building in Cold War computer science. In *Communities of Computing: Computer Science and Society in the ACM*. T.J. Misa, Ed. Morgan & Claypool, 2017, 199–227.
22. Wooldridge, M. *A Brief History of Artificial Intelligence: What It Is, Where We Are, and Where We Are Going*. Flatiron 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.

Copyright held by owner(s)/author(s).

^c See <http://bit.ly/409drns>

^d See <https://bit.ly/3PktK9l>