

Broken Promises & Empty Threats: The Evolution of AI in the USA, 1956–1996

Colin Garvey*

DOI: 10.15763/jou.ts.2018.03.16.02

Artificial Intelligence (AI) is once again a promising technology. The last time this happened was in the 1980s, and before that, the late 1950s through the early 1960s. In between, commentators often described AI as having fallen into “Winter,” a period of decline, pessimism, and low funding. Understanding the field’s more than six decades of history is difficult because most of our narratives about it have been written by AI insiders and developers themselves, most often from a narrowly American perspective.¹ In addition, the trials and errors of the early years are scarcely discussed in light of the current hype around AI, heightening the risk that past mistakes will be repeated. How can we make better sense of AI’s history and what might it tell us about the present moment?

* Copyright 2018 Colin Garvey. Colin Garvey is a Ph.D. candidate in the Department of Science and Technology Studies at Rensselaer Polytechnic Institute.

¹ Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence* (San Francisco: W. H. Freeman, 1979); Daniel Crevier, *AI: The Tumultuous History of the Search for Artificial Intelligence* (New York, NY: Basic Books, 1993); Rodney A. Brooks, *Cambrian Intelligence: The Early History of the New AI*, A Bradford Book (Cambridge: MIT Press, 1999); Nils Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Accomplishments* (New York: Springer, 2010).

This essay adopts a periodization used in the Japanese AI community to look at the history of AI in the USA. One developer, Yutaka Matsuo, claims we are now in the third AI boom.² I borrow this periodization because I think describing AI in terms of “booms” captures well the cyclical nature of AI history: the booms have always been followed by busts. In what follows I sketch the evolution of AI across the first two booms, covering a period of four decades from 1956 to 1996. In order to elucidate some of the dynamics of AI’s boom-and-bust cycle, I focus on the *promise* of AI. Specifically, we’ll be looking at the impact of statements about what AI one day would, or could, become.

Promises are what linguists call “illocutionary acts,” a kind of performance that commits the promise maker to a “future course of action.”³ A statement like, “We can make machines that play chess, I promise” has the potential to become true, if the promise is kept. But promises can also be broken. Nietzsche argued over a century ago that earning the right to make promises was a uniquely human problem. Building on that insight, the anthropologist Mike Fortun has explored the important role promises play in the construction of technoscience.⁴ AI is no exception. In Booms 1 and 2, the promises about AI were many, rarely kept, and still absolutely essential to its funding, development, and social impacts.

² 松尾豊(Yutaka Matsuo), 人工知能は人間を超えるか: ディープラーニングの先にあるもの. *Jinkō Chinō Wa Ningen o Koeru Ka: Dīpu Rāningu No Saki Ni Aru Mono*. Shohan. Kadokawa EPUB Sensho 021 (Tōkyō-to Chiyoda-ku: Kabushiki Kadokawa, 2015).

³ Antonio Blanco Salgueiro, “Promises, Threats, and the Foundations of Speech Act Theory,” *Pragmatics* 20:2 (2010): 213.

⁴ Mike Fortun, *Promising Genomics: Iceland and DeCODE Genetics in a World of Speculation* (Berkeley: University of California Press, 2008), 10.




Figure 1. The Three Booms of AI, an original diagram inspired by Yutaka Matsuo (see note 2).

The first boom was already being described as an era of “good, old-fashioned artificial intelligence” by the 1980s.⁵ Boom 1 was tied to the optimism that followed from the creation and multiplication of powerful electronic digital computers after the Second World War. Forged in military research centers and three major academic hubs (MIT, Carnegie-Mellon, and later Stanford University), this era was dominated by now-legendary AI pioneers such as Herbert Simon, Alan Newell, John McCarthy, and Marvin Minsky. They attempted to produce machines to perform certain tasks regarded as the purest form of intellectual activity by the elite, mostly white men, involved in shaping the trajectory of early computing—tasks like playing checkers and chess, as well as solving logic puzzles.

All the AI pioneers were consummate promise makers, but here we focus on Simon and Newell because their promises proved instrumental in shaping the field. In

⁵ Though the acronym “GOFAI” is used throughout the AI community to describe this period, the term originates with philosopher John Haugeland, in his *Artificial Intelligence: The Very Idea* (Cambridge: MIT Press, 1985).

1957, a year after the first conference on AI was held at Dartmouth College, the USSR launched *Sputnik*, the first orbital satellite, bringing Cold War tensions with the United States to new heights. In response, the American Department of Defense hastily formed the Advanced Research Projects Agency (ARPA, later DARPA) in 1958 to facilitate research and development of military and industrial technologies. Known for creating what would eventually become the Internet, ARPA was also the primary funder of AI R&D across the first two booms.⁶ AI's dependence on ARPA funding originated around the same time, when Simon and Newell presented a series of what they called "predictions" to a meeting of the Operations Research Society of America, a group with deep ties to the military.⁷ Describing the early work in AI, they confidently stated that "within ten years a digital computer" would become "the world's chess champion," "discover and prove an important new mathematical theorem," and even "write music."⁸ Simon and Newell summarized the status of Boom 1 by stating frankly that "*there are now in the world machines that think, that learn, and that create.*" They added that the machines' ability to do these things will "increase rapidly until—in a visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied."⁹ That is, the machines that will rival and ultimately overcome humans are already here.

⁶ Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America*, Inside Technology (Cambridge: MIT Press, 1996).

⁷ Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, 25th anniversary update (Natick: A.K. Peters, 2004).

⁸ Herbert A. Simon and Allen Newell, "Heuristic Problem Solving: The Next Advance in Operations Research," *Operations Research* 6, no. 1 (February 1958): 7–8.

⁹ Ibid.



Figure 2. Portrait of Herbert Simon by Richard Rappaport. Collection of Carnegie Mellon University: Herbert A. Simon. Creative Commons.

Inundated as we are by promissory technosciences, statements like Simon and Newell's would hardly raise an eyebrow today. But the scientific culture of the late-1950s was different, and their claims drew the ire of colleagues. The renowned US mathematician, Richard Bellman, was one of the first to take issue with their statements, noting with wry humor that while "magnificent," they were "not scientific writing."¹⁰ For Bellman, making such fantastic promises in a scientific journal was simply irresponsible. Simon and Newell replied that it was *because* they were such responsible scientists that they could make such promises: they were merely doing their duty by alerting the world to immanent developments in intelligent machinery.¹¹ Comparing themselves favorably to the irresponsible Nostradamus, Simon and Newell insisted their intention was to

¹⁰ Richard Bellman, "On 'Heuristic Problem Solving' by Simon and Newell," *Operations Research* 6, no. 3 (June 1958): 449.

¹¹ Herbert A. Simon and Allen Newell, "Reply: Heuristic Problem Solving," *Operations Research* 6, no. 3 (June 1958): 449.

“induce some intellectual leaders in our society to concern themselves with important developments that we judge to be only a decade off.”¹² This elaboration highlights the promissory dimension of their statements, which successfully induced the newly-formed ARPA to pay close attention to AI—and to begin funding Simon, Newell, and other AI pioneers at unprecedented levels for the next decade.¹³

These and other promises the AI pioneers made contributed directly to the perception that AI, as part of the ensemble of technologies then called “cybernation” (a portmanteau of “cybernetics” and “automation”), posed a considerable risk to the US labor force. President Kennedy referred to technological unemployment as the major domestic challenge of the 1960s.¹⁴ As an indication of the impact that these and other statements about machines “now in the world” were having on American society, the US Department of Labor’s Bureau of Labor Statistics’ two-volume annotated bibliography *Implications of Automation and Other Technological Developments* listed over 800 publications concerned with this theme between 1956-63.¹⁵ In 1964, the Ad-hoc Committee on the Triple Revolution petitioned President Johnson to act on the risks of “cybernation,” and later that year he launched the National Committee on Technology, Automation, and Economic Progress.¹⁶

Yet fears of “cybernation” were fading by the late 1960s as it became clear that the AI pioneers had made promises they could not keep. When a new decade opened and the magnificent machines “that think, that learn, that create”—that had been *promised*—were nowhere to be found, ARPA sobered up and began to cut AI funding.¹⁷

¹² Ibid.

¹³ Edwards, *The Closed World*, 261.

¹⁴ Robert Waterman McChesney and John Nichols, *People Get Ready: The Fight against a Jobless Economy and a Citizenless Democracy* (New York: Nation Books, 2016), 82.

¹⁵ Bureau of Labor Statistics, “Implications of Automation and Other Technological Developments: A Selected Annotated Bibliography” (US Department of Labor, December 1963).

¹⁶ McChesney and Nichols, *People Get Ready*, 82.

¹⁷ Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence* (San Francisco: W. H. Freeman, 1979); Edwards, *The Closed World*.

Remembering how “things had gotten a little bit weird” by the early 1970s, the prominent AI scientist Hans Moravec pointed to the role promising played in bringing about the funding cuts:

Many researchers were caught up in a web of increasing exaggeration. Their initial promises to D/ARPA had been much too optimistic. Of course, what they delivered stopped considerably short of that. *But they felt they couldn't in their next proposal promise less than in the first one, so they promised more.*¹⁸

(Author's emphasis.)

Moravec describes how ARPA, having been “duped,” decided that “‘some of these people were going to be taught a lesson,’ were having their two-million-dollar-a-year contracts cut to almost nothing!”¹⁹ Edward Feigenbaum, Simon's star student, remembered this as the “end of Eden and the beginning of the real world.”²⁰ AI was no longer seen as a promising technology, but as a risk. Consequently, the promises of AI scientists were no longer worth anything. AI lost credibility, and with it, the ability to make promises. Boom 1 went bust and AI entered the first Winter.

*

Feigenbaum had his own plan to revive AI. It consisted of two parts, technical and social. The technical part was a new kind of AI called “expert systems.” Abandoning his mentors' generalized approach to intelligence, Feigenbaum believed the key to intelligent machines lie in modeling the domain-specific knowledge of human experts. Rather than “general problem solvers,” AI should instead focus on building “expert systems” with the specific skills of professional chemists, doctors, and so on. The social part of his plan addressed funding. If the quantity of money and support the field had grown used to could not be procured with promises, then perhaps the *threat* of

¹⁸ Daniel Crevier, *AI: The Tumultuous History of the Search for Artificial Intelligence* (New York: Basic Books, 1993), 115.

¹⁹ *Ibid.*, 117.

²⁰ Nils Nilsson, Oral History of Edward Feigenbaum, June 20, 2007, CHM Reference number X3896.2007, Computer History Museum.

AI would do the trick. Feigenbaum intended to out-do his predecessors by securing attention and funding not only from the military, but industry and government more broadly. And as it happened, all those audiences were “quite interested, at the time, in the surge of Japan into the world’s arena of technology and business.”²¹ Fueled on exports of automobiles and consumer electronics, Japan’s late-1970s economic rise was increasingly the focus of the very elites whose support Feigenbaum needed to resuscitate AI. This growing concern about Japanese economic power presented an opportunity, but he would have to somehow connect AI to Japan. Luckily for him, around this time he received notice that the Japanese government was planning a large-scale AI project based loosely on his own expert systems.²² This Japanese project could be made to serve his purposes: If the military, ARPA, and other funders could no longer be sold on the promise of AI, perhaps they would buy into the *threat* that Japan would realize that promise.

Called “Fifth Generation Computer Systems” (FGCS), the Japanese project was a large-scale national computing project organized by the Ministry of International Trade and Industry (MITI), the government organization hailed for organizing Japan’s post-WWII “miracle economy.”²³ Scheduled to run for a decade in three phases, the FGCS project was funded at \$500 million and intended to produce both new computing hardware as well as AI software for the Japanese society of the 1990s.²⁴

²¹ *Ibid.*, 49–50.

²² *Ibid.*, 50.

²³ Chalmers Johnson, *MITI and the Japanese Miracle: The Growth of Industrial Policy, 1925-1975* (Stanford: Stanford University Press, 1982).

²⁴ JIPDEC, “Preliminary Report on Study and Research on Fifth-Generation Computers 1979-1980” (Japan Information Processing Development Center, Fall 1981).



Figure 3. The JW-10, the first Japanese word processor (Toshiba Science Museum, Kawasaki, Japan). Improving Japanese language computer interfaces was a major goal of the FGCS project.

Extensive planning for the FGCS began in the late-1970s, and it included consultations with Feigenbaum and other Western AI experts. When the project was announced at an international conference in Tokyo, 1981, Feigenbaum gave an invited keynote, praising the project's prospects and even calling for more international collaboration.²⁵ Yet immediately after returning to the United States, he began traveling to computer science departments around the country at his own expense to raise alarms about the threat the Japanese project posed.²⁶ Kicked out of Eden, starved of funding, and stripped of the right to make promises about AI, Feigenbaum could still make threats—and a threat is just as good as a promise. Both are statements about the future that commit the speaker to course of action.

This rough equivalence gave Feigenbaum exactly what needed to revive AI in the USA. He made his case with AI chronicler Pamela McCorduck in *The Fifth Generation: Japan's Computer Challenge to the World* (1983). Here they claimed that the FGCS project posed a significant threat to the United States for three reasons. First, computerization was bringing about a “New Wealth of Nations” in which data, information, and knowledge will be more important than material resources. Second,

²⁵ JIPDEC, “Proceedings of the International Conference on Fifth Generation Computer Systems, October 19-22, 1981,” FGCS International Conference (Japan Information Processing Development Center, 1981), <http://www.jipdec.or.jp/archives/publications/J0002118>.

²⁶ Nilsson, Oral History of Edward Feigenbaum.

Japan, lacking natural resources, had no choice but “to dominate the traditional forms of the computer industry” and “establish a ‘knowledge industry’.” Given these two premises, Japan’s launch of a national AI-focused computing program could clearly be nothing other than its attempt to take the global lead in the “knowledge industry” and “thereby become the dominant industrial power in the world.” Japan’s FGCS project thus threatened to rupture the global economic order, disrupt American hegemony, and consign “our nation to the role of the first great postindustrial agrarian society.”²⁷ Feigenbaum and McCorduck argued that in response, “Americans should mount a large-scale concentrated project of our own; that not only is it in the national interest to do so, *but it is essential to the national defense*”²⁸ (Author’s emphasis).



Japan, whose exports of cheap and reliable automobiles and consumer electronics propelled it to become arguably the most productive economy in the world by the late 1970s, came to be seen by many Americans as a more serious threat than Soviet nuclear weapons as “Japanophobia” gripped the USA for several years in the early 1980s. Photograph taken by the author.

The irony is that the FGCS, as the first of Japan’s many post-WWII national computing projects to be organized around basic research, was no such threat. The project did not include plans for the production of any commercial products, much less

²⁷ Edward A. Feigenbaum and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan’s Computer Challenge to the World* (Reading: Addison-Wesley, 1983), 1–3.

²⁸ *Ibid.*, 216.

for domination of the global economy. But a fair description of the project was not Feigenbaum and McCorduck's intention. Published at the height of "Japanophobia" in the USA, *Japan's Challenge* transmuted the launch of the Japanese AI project into a threat akin to the launch of *Sputnik* decades earlier—something both decision makers and the broader public could easily understand during the renewed Cold War tensions of the early 1980s.²⁹ Framed thusly, the Japanese "threat" was used to justify a wave of AI-focused national projects.

Within months, the United States responded with two major counter-threats: the Microelectronics and Computer Technology Corporation (MCC) and the Strategic Computing Initiative (SCI). The MCC was a for-profit consortium of 13 corporations. The *IEEE Spectrum's* special report on the new initiatives described it as the "most direct business counter" to the Japanese project.³⁰ The SCI, a billion-dollar, decade-long DARPA-funded computing project, quickly followed.³¹ Feigenbaum's plan to reinvigorate American AI was succeeding.³² So well, in fact, that an 'arms race' dynamic took hold of AI. For example, the massive SCI budget inspired an even larger program, the European Strategic Programme on Research in Information Technology (ESPRIT),

²⁹ For more on Japanophobia, see Anna Lowenhaupt Tsing, *The Mushroom at the End of the World: On the Possibility of Life in Capitalist Ruins*, 2015, Chapter 8. On the resemblance to *Sputnik*, Feigenbaum and McCorduck said America "needs a national plan of action, a kind of space shuttle program for the knowledge systems of the future." Feigenbaum and McCorduck, *The Fifth Generation*, 3.

³⁰ M. A. Fischetti, "The United States," *IEEE Spectrum* 20, no. 11 (November 1983): doi:10.1109/MSPEC.1983.6370021, 52.

³¹ Alex Roland and Philip Shiman, *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983-1993* (Cambridge: MIT Press, 2002).

³² When asked if his book had helped "get the so-called Strategic Computing Program [sic] started at DARPA?," Feigenbaum replied, "It had a huge influence in that. It got congressional staffs behind that. [...] The *Fifth Generation* book definitely helped sell that project." Nilsson, Oral History of Edward Feigenbaum, 51.

funded at \$1.5 billion over 5 years.³³ Along with the \$500 million Alvey project in the UK, these competing national computing projects fueled AI Boom 2.³⁴

The new AI boom was not without its critics, however. In his incendiary 1983 review of *Japan's Challenge*, AI scientist Joseph Weizenbaum, author of the famous ELIZA program, equated the *promise* of Boom 1 with the *threat* of Boom 2. After summarizing Feigenbaum and McCorduck's characterization of the threat posed by the FGCS project and the technologies it promised to produce, Weizenbaum agreed that the Japanese appear to be making "a very ambitious claim." He added that nearly identical claims have been made before—specifically by Simon and Newell in 1958. "[The] most recent ambitions of the Japanese," he explained, "were already close to being realized according to leaders of the American artificial intelligence community a quarter of a century ago!"³⁵

However, empty threats are only as valuable as promises that cannot be kept. By the late 1980s, AI Boom 2 was already faltering. AI critics Hubert and Stuart Dreyfus could remark with relish, "now that the *promised* expert systems have been marketed and used, reality, buried under years of dreams and distortions, has begun to assert itself."³⁶ Comparing the promissory rhetoric of the early 1980s to dire assessments from the second half of the decade, they argued that another Winter is imminent. Responsibility for the coming crash, however, belongs with the AI scientists themselves, as "it seems unfair to blame the managers for failing to sell a product that did not perform as *promised*."³⁷

³³ H. Nasko, "European Common Market," *IEEE Spectrum* 20, no. 11 (November 1983): 71–72, doi:10.1109/MSPEC.1983.6370023.

³⁴ Brian Oakley and Kenneth Owen, *Alvey: Britain's Strategic Computing Initiative* (Cambridge: MIT Press, 1989).

³⁵ Joseph Weizenbaum, "The Computer in Your Future," *The New York Review of Books*, October 27, 1983, <http://www.nybooks.com/articles/1983/10/27/the-computer-in-your-future/>.

³⁶ Hubert L Dreyfus and Stuart E Dreyfus, *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer* (New York: Free Press, 1986), ix.

³⁷ *Ibid.*, xi.

Boom 2 came crashing down as the 1980s closed, and Japan did not dominate the globe with AI. Instead, its economy collapsed in the early 1990s. The Japanese FGCS was consequently described as a failure by the US computing community, and many rewrote history to claim they had always seen the FGCS as an empty threat.³⁸ None bothered to hold the counter-threats—MCC, SCI, Alvey, and ESPRIT—to similar account. Rather, a representative statement by an attendee at an SCI conference summarizes the Machiavellian way all of these Western projects would be evaluated in their home countries: “from my point of view, [SCI] is already a success. We’ve got the money.”³⁹

AI’s lowest point came in a scarcely-remembered 1996 confrontation between top-ranked chess grandmaster Gary Kasparov and IBM’s Deep Blue, an advanced expert system built specifically to play chess. Kasparov handily defeated the machine, four games to two. Thus, as of the mid-1990s, AI was little more than a morass of broken promises and empty threats, a defeated technoscience impaled on its own hubris. This second AI Winter would only begin to thaw in 1997, when Deep Blue defeated Kasparov in a highly publicized and highly controversial re-match.⁴⁰ With this event, AI finally realized the promise Simon and Newell made on its behalf forty years prior.

Boom 3 would not begin for more than another decade, however, when the revival of neural network-based algorithms unleashed “machine learning” as the

³⁸ Ehud Shapiro and David H. D. Warren, “Personal Perspectives,” *Communications of the ACM* 36, no. 3 (March 1993): doi:10.1145/153520.153534, 46–49; McCorduck, *Machines Who Think*, 2004; Nils Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Accomplishments* (New York: Springer, 2010). Compare for example the hysterical tone of David H. Brandin, “ACM President’s Letter: The Challenge of the Fifth Generation,” *Communications of the ACM* 25, no. 8 (August 1982): 509–10, with the sober assessment found in David H. Brandin and Michael A. Harrison, *The Technology War: A Case for Competitiveness* (New York: Wiley, 1987).

³⁹ Quoted in Roland and Shiman, *Strategic Computing*, 94.

⁴⁰ G. K. Kasparov and Mig Greengard, *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins* (New York: PublicAffairs, 2017).

successor to “expert systems.” Relying on vast quantities of data and enormous computing power instead of the knowledge of human experts, machine learning has led to several recent breakthroughs in areas like image recognition and in games such as Go and poker. Advocates are saying the “amazing AI we were promised is coming, finally”.⁴¹ Yet the promises are already turning to threats as an arms race dynamic takes hold, this time between the United States and China.⁴² Unlike the Japanese “threat” that animated Boom 2 however, the threat now driving Boom 3 may have some substance behind it. China’s recent “Next Generation Artificial Intelligence Development Plan” explicitly states the nation’s goal of world domination in AI by 2030, through massive investments in R&D and a policy of “military-civil fusion.”⁴³ Only time will tell, however, if the Chinese AI community can keep that promise.

Suggested Readings

Hubert L. Dreyfus, *What Computers Still Can’t Do*. New York: MIT Press, 1992.

Tessa Morris-Suzuki, *The Technological Transformation of Japan: From the Seventeenth to the Twenty-First Century*. New York: Cambridge University Press, 1994.

Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge: MIT Press, 1996.

⁴¹ Vivek Wadhwa, “The Amazing Artificial Intelligence We Were Promised Is Coming, Finally,” *Washington Post*, June 17, 2016, <https://www.washingtonpost.com/news/innovations/wp/2016/06/17/the-amazing-artificial-intelligence-we-were-promised-is-coming-finally/>.

⁴² Elsa Kania, “The Next U.S.-China Arms Race: Artificial Intelligence?” *The National Interest*, March 9, 2017, <http://nationalinterest.org/feature/the-next-us-china-arms-race-artificial-intelligence-19729>; Julian E. Barnes and Josh Chin, “The New Arms Race in AI,” *Wall Street Journal*, March 2, 2018, <https://www.wsj.com/articles/the-new-arms-race-in-ai-1520009261>.

⁴³ Gregory Allen and Elsa B. Kania, “China Is Using America’s Own Plan to Dominate the Future of Artificial Intelligence,” *Foreign Policy*, <https://foreignpolicy.com/2017/09/08/china-is-using-americas-own-plan-to-dominate-the-future-of-artificial-intelligence/> (accessed September 10, 2017).

Alex Roland and Philip Shiman, *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983-1993*. Cambridge: MIT Press, 2002.

Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. Natick: A.K. Peters, 2004.