Implications for AI Research: Applying Lessons from the Expert Systems Boom and Bust to the Current Large-Language Model Boom

Ted E. Senator

Abstract

AI is currently in the midst of a boom, mostly due to the success and predominance of large language models and associated models for other perceptual tasks such as computer vision. Yet AI has experienced several booms and busts over the past 75 years. While the booms are driven by commercial potential, the following busts affect not only commercial investment but also research funding and trends. This paper examines the expert systems boom of the 1980s and the following AI Winter, identifies similarities, analogs, and differences with the current boom, and projects potential outcomes and directions for AI research that may follow when the current enthusiasm wanes based on these similarities, analogs, and differences. The presentation is distinct from currently active technical discussions and debates about the potential and limitations of large models such as whether problems of hallucination will be solved, whether they can reason, or whether they will achieve AGI; rather, it examines previous AI techniques and how they evolved once their capabilities and limitations became well understood.

Introduction

Large Language Models (LLMs) are currently fueling a boom, some might say a mania, for AI. While there is extensive discussion in the AI community about the potential and limitations of LLMs, there is minimal awareness of prior AI booms and the follow-on busts, or AI Winters. This paper reviews the history of prior AI booms, with particular focus on the Expert Systems boom of the 1980s and their aftermath and maps this history to current and potential future paths for AI research in a post-LLM environment.

An Extremely Abbreviated History of AI

AI as a discipline arose at the Dartmouth Workshop in 1956. Initial approaches were based on logic; this led to the first AI Winter when logic could not scale to or represent the richness of real-world problems. Machine translation was deemed a failure in the ALPAC Report (Pierce, Carroll et al. 1966). (An alternative approach to logic was based on perceptrons, but this was largely abandoned when Minsky and Papert in 1969 showed that they were incapable of computing the XOR and other predicates.) Marked by the 1973 Lighthill Report (Lighthill 1973), the first AI Winter occurred in the mid-1970s through 1980. It was replaced by renewed optimism characterized by the promise of expert systems and the rise of specialized Lisp machines. The phrase "in the knowledge lies the power" described this age. Expert System Shell companies were started. The US Strategic Computing Initiative and Japans's Fifth Generation Project were major government funded research efforts. This age ended when it became clear that expert systems were brittle, expensive to develop and maintain, and not easily adaptable. Despite research into knowledge acquisition and even practical guides to this art (Scott, Clayton, and Gibson 1991), the knowledge-engineering bottleneck was believed to be the major impediment to achieving special purpose AI. And the specialized workstations were overtaken by general purpose machines at far lower costs.

The second AI Winter in the late 1980s to mid-1990s was a period when case-based reasoning was popular; this topic dominated AAAI-90. Research funding tended to favor specific applications such as image understanding, natural language processing, planning and reasoning, etc., rather than more general approaches to AI. Statistical learning began to emerge from research labs, and the hope was that knowledge could be learned rather than hand-coded. IBMs Deep Blue commanded popular attention when it defeated the world chess champion Gary Kasparov in 1997. Much attention was focused on the internet, and the AI community was not immune. Google revolutionized web search with its PageRank algorithm in 1998.

The 2000's were the era of big data. The terms "data mining" and "knowledge discovery from data" were introduced, relying on techniques from machine learning, with an annual conference that began in 1995 sponsored by AAAI and spun-out of in 1999. Google, Facebook, and Amazon used search and recommendation techniques originally developed in the AI community. Specialized applications in many areas such as fraud detection, telecommunications networks and management, medicine, spacecraft operations, diagnosis, customer support, and many others began to appear. Most prominently, IBM's Watson won the Jeopardy! challenge in 2011.

The major breakthrough in 2012 in computer vision and the development of transformers in 2017 enabled current deep learning based approaches. So-called Foundation Models, specialized graphics processing unit chips, and the availability of internet-scale training data have led to major advances. Widespread commercial interest and the extensive infrastructure requirements of these large models have led to a movement of researchers to industry from academia. Societal questions such as how to trust these models and whether they will lead to AGI and eventually overtake humans as a dominant species are debated both in AI circles and popular media. LLMs dominate research conferences, Chatbots are in widespread use, and other significant applications based on foundation models such as LLMs are beginning to appear. And newer techniques such as Agentic AI are emerging rapidly.

The Expert Systems Boom

Expert Systems were originally conceived to use the knowledge of human experts in specific domains. They represented knowledge as production rules, and used forward or backward chaining to reason about specific problems. A key aspect of expert systems was the separation of the explicit domain knowledge (i.e., the rules) from the inference engine. The term knowledge based systems is a generalization of expert systems; it removes the goal of performing at the level of human experts but still is capable of using knowledge to reason at the journeyman level and to distribute knowledge consistently across organizations. Notable examples that led to the widespread research and commercial interest in expert systems were PROSPECTOR, which was used to identify deposits of valuable metals, R1, that configured minicomputers sold by the Digital Equipment Corporation, DENDRAL, which performed chemical structure elucidation, and CADUCEUS and MYCIN, both of which used internal medicine knowledge to diagnose complex sets of symptoms.

Expert systems were a major focus of AI research and applications in the 1980s. Early AAAI conferences, while much smaller than today's, had papers with Expert System in their titles. The first AAAI conference in 1980 had 8 papers with Expert or Expert System in their titles; these included both reports of applications and of research. There were also several papers with titles describing applications that used Expert System techniques and others that used the more general term "Knowledge-Based System." A characteristic of the AAAI conference during the early 1980s is that both research and application papers were included; many papers reported on research advances in the context of applications, although not at the scale needed for deployment. AAAI-82 had 9 nine papers with the word Expert or Expert System in the title, but 23 that described applications of which only four used the word "expert" in the title; AAAI-83 19 applications in a category called Expert Systems (of which 6 had the word "Expert" in the title); and AAAI-84, six in the category with only two explicitly called Expert Systems. AAAI-86 recognized the distinction between science and engineering and had two corresponding tracks: 18 papers in the engineering track were classified as applications; the remainder were various enabling technologies. 1986 was the first year that the term "knowledge acquisition" appeared with four papers in this category. AAAI-1987 saw a grouping of machine learning and knowledge acquisition, recognizing perhaps that manual acquisition of knowledge needed to be enhanced for effective construction of applications. Papers this year were grouped by task domain, such as reasoning, planning, and natural language; there was no explicit applications category [1].

In 1989 the Innovative Applications of AI (IAAI) conference was established to complement the research oriented AAAI conference. Only deployed applications were eligible for inclusion. Of the 30 papers in 1989, six titles used the word "expert" and six "knowledge"; in 1990 the corresponding numbers were seven and three out of 22; in 1991 eight and four of 21; in 1992 one and seven of 19; in 1993 five and six of 16; in 1994 five and zero of 16; in 1995 only one and three of 17; and in 1996 four and one of 18. In 1997 IAAI expanded to include emerging as well as deployed applications so the numbers are no longer comparable. During this period of IAAI, expert systems no longer appeared in the titles of papers of the AAAI conference. (Case based reasoning became very popular during this time as a newer technique that could potentially overcome some of the shortcomings of expert systems.)

In parallel with the AI research community's emphasis on expert systems, the 1980s were a period of great commercial interest. Estimates of the investment are typically around $1B annually in 1985 alone. Major firms such as IBM, Digital Equipment, Texas Instruments, General Electric and General Motors developed systems for internal use. Companies specializing in knowledge engineering, the practice of building expert systems, were formed, the most noticeable being Teknowledge, Inference Corp., Carnegie Group, and Intellicorp. They provided consulting services and tools, the latter called "expert system shells." These tools turned out not to be a viable business, as the market demanded vertical solutions populated with domain-specific knowledge rather than empty generic shells, because the major cost for developing an expert system was acquiring and representing the domain knowledge. None of these companies exist today: Teknowledge spun off its government services business as ISX Corporation, later acquired by Lockheed Martin, and the original Teknowledge ceased to operate in 2007-2010. Inference Corporation was acquired in 1999. Its Automated Reasoning Technology (ART), itself a derivative of OPS5 — the forward-chaining rule-based production system used in an often cited example of an early expert system success, the Digital Equipment Corporation's R1/XCON system to configure their minicomputers — influenced NASA's CLIPS in the mid-1980s. Carnegie Group was acquired by the UK firm Logica PLC in 1998. And Intellicorp, with their backward-chaining Knowledge Engineering Environment (KEE) tool, converted KEE to run on PCs and sold its assets to Tricentis in 2019.

The 1980s also saw the introduction of specialized workstations on which to run expert systems software. Most notable are Symbolics, Lisp Machine Incorporated, and Texas Instruments. These machines not only were optimized to ex-

[1] There was no AAAI in 1981 or 1985 due to the practice of not having this conference when IJCAI was in North America

ecute rule-based processing of expert systems — the rule-based processing software was typically written in LISP, the language used for most AI research — but often came with sophisticated software engineering environments necessary for AI programming. However, with the rise of UNIX-based general purpose workstations from companies such as Sun and Apollo and their much larger customer bases, these specialized AI machines could not keep up and were superseded. Lisp was also superseded as the basis for rule-based and other AI processing by languages such as C, C++, Java, and Smalltalk as evidenced by a set of textbooks authored by Patrick Henry Winston of MIT. Xerox also offered their version of LISP called Interlisp, optimized for their proprietary hardware.

The Expert System boom was characterized by excitement among researchers and investors. Articles such as (Prerau 1985) and (Bobrow, Mittal, and Stefik 1986) had as a subtext the idea that qualified AI scientists were a scarce resource and we would have the luxury of choosing which projects to work on from many potential sponsors. While this attitude may have been profitable in the short-term, it likely discouraged many potential future sponsors.

As people began to understand the key limitations of expert systems such as their brittleness (i.e., when processing something at the boundary or outside the scope of their knowledge their results were unreliable at best), their lack of common sense, their inability to learn (necessary because the world is dynamic), their difficulty in handling incomplete or uncertain information, and the expense involved not only in developing but also maintaining, they were viewed as a failed technology. AI projects also failed for non-technical reasons, just like other software projects, as noted in (Jacobstein 2008). Finally, the problems addressed by expert systems were often only a small part of what a professional did, so minimizing the effort had little effect on overall productivity. For example, solving the medical diagnosis problem didn't really benefit doctors, because diagnosis is a small part of their job. And in large enterprises in which multiple human agents process only parts of a set of transactions (e.g., different agents may touch an airline reservation over multiple days), the core technology is a database rather than AI, so the AI must integrate closely with the database management system.

However, there were positive aspects too. In particular, new research directions in machine learning (to construct knowledge bases automatically), alternative representations of knowledge (most especially case-based reasoning, which became popular around 1990 — although applications often comprise only retrieval of similar cases and ignore the aspect of adaptation that was also essential in the case-based reasoning paradigm — as a way of capturing past experience and expertise), and continued research on the fundamentals of knowledge representation and reasoning. An often overlooked benefit was the systematic formalization of knowledge in specific domains. Perhaps most significant was the CYC project (Lenat 1995), that attempted to build a repository of formally represented common-sense knowledge that could be used in many domains. And some of the techniques pioneered by AI researchers during the expert systems boom influenced other directions in computer science such as object-oriented programming languages.

How Expert Systems Evolved

What happened with the technologies of expert systems? The short answer is that they were incorporated into other technologies that were more widespread and fundamental. For example, rule-based processing appeared as a feature of Oracle and other database management systems. The emphasis on stand-alone AI systems made this path a dead-end; as expert systems were integrated with other corporate systems their features needed to be incorporated into the standard IT infrastructure, not only the hardware but also the large-scale databases that supported transaction processing as the major use of IT in large enterprises. Expert system shells became available on PCs as well as workstations and can be used to complement other technique in enterprise-scale applications.

What lessons can we learn from how expert systems evolved?

First, the most obvious lesson is that the AI research community should be careful about our claims. Letting the excitement about new possibilities override a careful analysis of strengths, weaknesses, benefits and limitations of a new technology resulted in an AI Winter, during which funding for research as well as commercial investment were significantly reduced. It is incumbent on us to be measured in our assessments of potential even when we are excited by major research breakthroughs. Communicating these subtleties to business and Government can be difficult however.

Second, while an AI technology can have enormous potential, it is often most useful as a component of a larger application. An AI technology that is viewed as the next breakthrough eventually becomes just another technique in the toolbox of techniques available to system designers and application developers, to be employed when appropriate and effective and to be integrated with other techniques that are more effective for other aspects of a problem.

Third, technology selection needs to be driven by business problems and benefits (i.e., demand) rather than by technologies in search of problems, as was done with expert systems during the height of the boom. Developers of AI applications need not only to have a thorough understanding of the business problems they are trying to solve, but also to be able to communicate clearly with the users.

The Current LLM Boom

Much like expert systems were key to the AI boom in the 1980s, Large Language Models are the main emphasis of today's AI research and applications. These large models can be thought of as another method to capture knowledge, except instead of it being hand coded (as in expert systems) or learned from data (as in the 1990s and 2000s), it is learned from the repository of natural language available on the internet. This knowledge is not curated and may be inaccurate, but it is extensive and easily accessible. The distributed representations as large numbers of weights or parameters makes it difficult if not impossible to examine by humans, however.

To quantify the current boom in Large Language Models in AI research, we have examined recent AAAI-Proceedings and counted paper titles that mention either the phrase "Large Language Model" or contain the abbreviation LLM. The AAAI-25 Proceedings contain 264 papers whose titles contain the phrase "Large Language Model" or the abbreviation LLM. AAAI-2024 has 72 papers with these titles, AAAI-2023 only two, and AAAI-2022 only one. Prior conferences have none [2]. As one might expect, applications lag research, and reports of deployed applications that mention LLMs appear for the first time only in IAAI-2025, with 4 of the 9 deployed application papers mentioning LLMs in their title.

Not only does the amount of research interest into Large Models dwarf the amount during the heyday of expert systems, but the commercial investments are even greater. According to the Stanford AI Index, corporate AI investment in 2024 exceeded $252B. The percentage of organizations using AI for at least one function has grown to 78%. However, according to the MIT NANDA project, 95% of generative pilots at companies are failing.

Like in the expert system days, specialized hardware is needed for training and execution of large AI models. However, unlike in the expert system days when this hardware was being produced in small amounts and purchased primarily by research labs, it is now coming from some of the most well capitalized corporations in existence, the most valuable of which is Nvidia, and is integrated with other computing infrastructure.

Models of Booms and What Happens Next

In light of the current boom, it is worthwhile to review the financial literature regarding how bubbles burst. Bubbles typically go through five stages according to the economist Hyman Minsky: (1) Displacement, (2) Boom, (3) Euphoria, (4) Profit-Taking, and (5) Panic. While this analysis is focused on financial assets, it can occur with respect to technologies as well. A more applicable model might be the well known Gartner Hype Cycle, in which the stages are (1) Innovation Trigger, (2) Peak of Inflated Expectations, (3) Trough of Disillusionment, (4) Slope of Enlightenment, and (5) Plateau of Productivity. As of this writing, Generative AI — the most recent application of LLMs, is just approaching the peak.

I propose a similar model as follows:

There are three ways a boom can end: (1) disillusionment, (2) realism, and (3) explosion.

Disillusionment is what occurs when the boom is due to unrealistic inflated expectations that are inevitably not met. Disillusionment resulting from AI bubbles has resulted in what has been called "AI Winters" — periods during which research funding dries up, proposals avoid the term AI (but sometimes invent synonyms, e.g., machine intelligence or automated cognition or use related terms, such as decision support, optimization, etc.), industry investment is eliminated, pilot projects are shut down, and articles about why the technology failed appear. This seems to characterize the path taken by Expert Systems.

Realism is what happens when a bubble, rather than bursting, deflates slowly until it reaches an equilibrium. This is of course the best-case outcome. In this case the technology's capabilities are thoroughly understood and its limitations recognized. Use cases where it can be successful are identified. It is treated not as a full-system solution but rather as a component that be be part of a larger system. This path corresponds to reaching Gartner's Plateau of Productivity without falling into the intervening Trough of Disillusionment. (Economists may call this a soft landing.)

Explosion is what happens when a boom is actually justified — when expectations are met and perhaps exceeded and the world is transformed by the new technology. The internet reached this stage, although with growing pains after the hype of the late 1990s. It is not inconceivable that LLM-based AI will result in this outcome.

Possible Paths Forward

There are three schools of thought regarding paths forward for AI, two extreme and one more realistic and I think — and hope — far more likely. First, we can not rule out that the goal of Artificial General Intelligence (AGI) will be achieved by scaling up current techniques to larger and larger models. This view is promoted by (Bubeck et al. 2023), although it is viewed more skeptically with the very recent release of GPT-5. The danger of this outcome is of course that of a super-human intelligence that becomes a threat to humanity. A second extreme outcome is to recognize the "arms race" among major AI companies, including new ones such as OpenAI, Anthropic and more traditional technology providers such as Alphabet/Google, META/-Facebook, Amazon, Apple, Microsoft, and others. Because of the arms race aspect and the perceived first-mover advantage, it could happen that large-model based applications are deployed without a careful understanding of the risks. This likely would lead to a disaster with major loss of money and/or life, and if this were to occur, would almost certainly lead to onerous and overly-broad Government regulations and a drying up of funding for both applications and for research. Another danger is that an AI could be misused by nefarious human actors to create sophisticated attacks or to mislead humans and create social unrest. And the massive energy requirements for large model based AI could have extremely negative environmental impacts. A third path is more hopeful: one of responsible AI, in which a thorough understanding of the capabilities and limitations of these new techniques develops, the nascent field of ethical, trustworthy AI continues to develop, the bubble deflates slowly, large models become just another tool in the AI developers toolbox, and new ideas take over the research community.

What might these new ideas be? There have been two recent complementary efforts to identify future directions for AI research: The AAAI Presidential Report on the Future of AI Research (Rossi et al. 2025) was published in March

[2]These numbers were obtained by searching the AAAI Proceedings from the relevant years and manually removing duplicates. Counts may be off in either direction by a small number, but these minor errors do not affect the general conclusions.

2025, and the CCC Whitepaper Envisioning Possible Futures for AI Research (Jensen et al. 2025) was published in July 2025. The AAAI Presidential Report was prepared by 25 experienced AI researchers with 15 additional contributors and 475 respondents from a community survey. It comprises 17 topics, some of which are technical and reflect current research trends (e.g., AI Reasoning, AI Agents, Embodied AI, etc.), some of which are application areas (e.g., AI for Social Good, AI for Scientific Discovery, etc.), and some of which are issues related to widespread adoption and use of AI (e.g., AI Ethics, AI Perception vs Reality, etc.) The CCC Whitepaper discusses six "AI Research Futures": Neuro-Symbolic AI, Neuromorphic AI, Embodied AI, Multi-Agent AI, Human-Centered AI, and Quantum AI. All of these areas are of active concern to the AI community and related research communities and are possible areas of new ideas and approaches as well as fruitful domains for AI research advances and breakthroughs, in the context of any of the three resolutions to the current AI bubble.

References

Bobrow, D. G.; Mittal, S.; and Stefik, M. J. 1986. Expert Systems: Peril and Promise. *Communications of the ACM*, 29(9): 880–894.

Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; and Zhang, Y. 2023. Sparks of Artificial General Intelligence: Early Experiments with GPT-4. *arXiv preprint*.

Defense Advanced Research Projects Agency (DARPA). 1983. Strategic Computing: New-Generation Computing Technology: A Strategic Plan for the Development and Application to Critical Problems in Defense. Technical report, DARPA. Available at https://apps.dtic.mil/sti/tr/pdf/ADA141982.pdf.

Hayes-Roth, F.; Waterman, D. A.; and Lenat, D., eds. 1983. *Building Expert Systems*. Reading, MA: Addison-Wesley.

Jacobstein, N. 2008. Often, It's Not About the AI. *AI Magazine*, 29(2): 54.

Jensen, D.; Danks, D.; Elbaum, S.; Regli, W.; Turk, M.; Wierman, A.; Yanco, H.; Maher, M. L.; and Griffin, H. 2025. Envisioning Possible Futures for AI Research. Technical report, Computing Research Association (CRA), Washington, D.C. Available online.

Lenat, D. B. 1995. CYC: A Large Scale Investment in Knowledge Infrastructure. *Communications of the ACM*, 38(11): 33–38.

Lighthill, J. 1973. Artificial Intelligence: A General Survey. In *Artificial Intelligence: A Paper Symposium*. UK: Science Research Council.

Minsky, M.; and Papert, S. A. 1969. *Perceptrons: An Introduction to Computational Geometry*. Cambridge, MA: MIT Press.

Pierce, J. R.; Carroll, J. B.; et al. 1966. Language and Machines: Computers in Translation and Linguistics. Technical report, National Academy of Sciences, National Research Council, Washington, D.C. ALPAC Report.

Prerau, D. S. 1985. Selection of an Appropriate Domain for an Expert System. *AI Magazine*, 6(2): 26.

Rossi, F.; Bessiere, C.; Biswas, J.; Brooks, R.; Conitzer, V.; Dietterich, T. G.; Dignum, V.; Etzioni, O.; Forbus, K. D.; Freuder, E.; Gil, Y.; Hoos, H.; Horvitz, E.; Kambhampati, S.; Kautz, H.; Kim, J.; Kitano, H.; Mackworth, A.; Myers, K.; De Raedt, L.; Russell, S.; Selman, B.; Stone, P.; Tambe, M.; and Wooldridge, M. 2025. AAAI 2025 Presidential Panel on the Future of AI Research. Technical report, Association for the Advancement of Artificial Intelligence (AAAI).

Scott, A. C.; Clayton, J.; and Gibson, E. 1991. *A Practical Guide to Knowledge Acquisition*. Reading, MA: Addison-Wesley.

Stanford Institute for Human-Centered Artificial Intelligence (HAI). 2025. AI Index Report 2025. https://hai.stanford.edu/ai-index/2025-ai-index-report. Accessed 2 September 2025.

Steele, G. L. 1984. *Common LISP: The Language*. Bedford, MA: Digital Press.

Winston, P. H. 1994a. *On To C*. Reading, MA: Addison-Wesley.

Winston, P. H. 1994b. *On To C++*. Reading, MA: Addison-Wesley.

Winston, P. H. 1997. *On To Smalltalk*. Reading, MA: Addison-Wesley.

Winston, P. H.; and Horn, B. K. P. 1981. *LISP*. Reading, MA: Addison-Wesley.

Winston, P. H.; and Narasimhan, S. 1996. *On To Java*. Reading, MA: Addison-Wesley.

Winston, P. H.; and Prendergast, K. A., eds. 1984. *The AI Business: Commercial Uses of Artificial Intelligence*. Cambridge, MA: MIT Press.