

The AI Bubble and the U.S. Economy: How Long Do ‘Hallucinations’ Last?

Servaas Storm *

Working Paper No. 240

September 23rd, 2025

ABSTRACT

The U.S. is undergoing an extraordinary, AI-fueled economic boom: The stock market is soaring thanks to exceptionally high valuations of AI-related tech firms, which are fueling economic growth by the hundreds of billions of U.S. dollars they are spending on data centers and other AI infrastructure. The AI investment boom is based on the belief that AI will make workers and firms significantly more productive, which will in turn boost corporate profits to unprecedented levels. But evidence is piling up that generative AI (GenAI) is failing to deliver. This paper argues that (i) we have reached “peak GenAI” in terms of current Large Language Models (LLMs); scaling (building more data centers and using more chips) will not take us further to the goal of “Artificial General Intelligence” (AGI); returns are diminishing rapidly; (ii) the AI-LLM industry and the larger U.S. economy are experiencing a speculative bubble, which is about to burst (because of the

* Department Economics of Technology and Innovation (ETI), Faculty of Technology, Policy and Management, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands.
S.T.H.Storm@tudelft.nl

The author is grateful to an unknown reviewer and Thomas Ferguson for insightful comments on and valuable improvements to the paper. The author accepts the blame for whatever deficiencies remain.

first point); and (iii) the U.S. bet the farm on a future dominated by U.S.-owned AGI, because, for geopolitical reasons, it could not afford to risk to lose the AI-race with China; this geopolitical bet on AGI is now going bad.

<https://doi.org/10.36687/inetwp240>

JEL codes: E24, F52, G01, O30, O33.

Keywords: Artificial intelligence; generative AI; AI bubble; ChatGPT; LLMs; productivity impacts; profitability; price-earnings ratio; scaling; hallucinations; energy and water use; geopolitics of AI race.

The infectious GenAI fever

The summer of 2025 did not bring good news for [enthusiasts of generative Artificial Intelligence](#) (GenAI) who were all hyped up by the inflated promise of the likes of OpenAI's Sam Altman that "Artificial General Intelligence" (AGI), the holy grail of current AI research, would be right around the corner. Already in January 2025, the honorable Mr. [Altman wrote](#) that "we are now confident we know how to build AGI". Altman's optimism echoed claims by OpenAI's partner and major financial backer [Microsoft, which had put out a paper in 2023](#) claiming that the GPT-4 model already exhibited "sparks of AGI." [Elon Musk](#) (in 2024) was equally confident that the Grok model developed by his company xAI would reach AGI, an intelligence "smarter than the smartest human being", probably by 2025 or at least by 2026. [Meta CEO Mark Zuckerberg](#) said that his company was committed to "building full general intelligence", and that [super-intelligence is now "in sight"](#). Likewise, [Dario Amodei](#), co-founder and CEO of Anthropic, said "powerful AI", *i.e.*, smarter than a Nobel Prize winner in any field, could come as early as 2026, and usher in a new age of health and abundance — the U.S. would become a "country of geniuses in a datacenter", if AI didn't wind up killing us all. And in a TED talk in April 2025, [Eric Schmidt](#), *soi-disant* tech guru and former Google CEO, stated that "the AI revolution is underhyped", predicting that we will have AGI "within three to five years."

For Mr. Musk and his GenAI co-travelers, the biggest hurdle on the road to AGI is the lack of computing power (installed in data centers) to train AI bots, which, in turn, is due to a lack of sufficiently advanced computer chips. According to [Musk](#), training the Grok 2 model took about 20,000 Nvidia H100 AI Graphics Processing Units (GPUs), while the Grok 3 model required 100,000 H100 chips. This information (dating from 2024) is already obsolete: industry sources expect Musk's xAI to be using between 500,000 and 1 million H100 equivalents in 2025. Note that xAI is a small player: Amazon, Google, Meta and Microsoft, taken together, are expected to be employing [more than 8 million H100 AI chips in 2025](#). [Nvidia's production](#) of these GPUs is estimated to reach 6.5 to 7 million units in 2025. The demand for more data and more data-crunching capabilities will require [about \\$3 trillion in capital](#) just by 2028, in the estimation of Morgan Stanley. [That would exceed the capacity of the global credit and derivative securities markets](#). Spurred by the imperative to win the AI-race with China, the GenAI propagandists firmly believe that the U.S. can be put on the yellow brick road to the Emerald City of AGI by building more data centers faster (an unmistakably "accelerationist" expression).¹

Interestingly, AGI is an ill-defined notion, and perhaps more of a marketing concept used by AI promoters to persuade their financiers to invest in their endeavors. According to ChatGPT, AGI "refers to a type of AI that can understand, learn, and apply knowledge across a wide variety of tasks at the same level or even better than humans. AGI would have the ability to perform any intellectual task that a human can."

¹ The Pentagon is investing \$800 billion in the development of *specific* military AI capabilities by Anthropic, Google, OpenAI and xAI in military AI ([here](#)). The U.S. became the largest investor in the [global spyware market in 2024](#) which heavily relies on AI tools. As pointed out by Thomas Ferguson, there is a growing overlap between firms and investors in high tech, defense and finance — and the political system (see [French 2025](#)).

In other words, an AGI model can generalize beyond specific examples found in its training data, similar to how some human beings can do almost any kind of work after having been shown few examples of how to do a task, by learning from experience and changing methods when needed. AGI bots will be capable of outsmarting human beings, creating new scientific ideas, and doing innovative as well as all of routine coding. AI bots will be telling us how to develop new medicines to cure cancer, fix global warming, drive our cars and grow our genetically modified crops. Hence, in a radical bout of creative destruction, AGI would transform not just the economy and the workplace, but also systems of health care, energy, agriculture, communications, entertainment, transportation, R&D, innovation and science.

OpenAI's [Altman boasted](#) that AGI can “discover new science,” because “I think we’ve cracked reasoning in the models,” adding that “we’ve a long way to go.” He “think[s] we know what to do,” saying that OpenAI’s o3 model “is already pretty smart,” and that he’s heard people say “wow, this is like a good PhD.” Announcing the launch of ChatGPT-5 in August, Mr. Altman posted on the internet that “We think you will love using GPT-5 much more than any previous AI. It is useful, it is smart, it is fast [and] intuitive. With GPT-5 now, it's like talking to an expert — a legitimate PhD level expert in anything any area you need on demand, they can help you with whatever your goals are.”

A good, old cold douche (per Joseph Alois [Schumpeter](#))

But then things began to fall apart, and rather quickly so. In non-chronological order, there was the letdown of the newly launched ChatGPT-5 relative to hyped-up expectations. Then came the shock brought about by three widely publicized, sobering research papers showing that novel ever-larger GenAI models do not become better, but worse, and do not reason, but rather parrot reasoning-like text. And, finally, markets got thoroughly spooked by an MIT study showing that 95% of GenAI pilots at corporations are falling flat. The future used to be better — not even very long ago. Before proceeding, it must be emphasized that this paper is not arguing against the usefulness of decentralized domain-specific AI applications that are already effective in specialized fields² — even if the cost of developing these applications has to be weighed against using the money for alternative (societal and scientific) purposes, for instance, clean energy and public health.

ChatGPT-5 is a letdown

The first piece of bad news is that much-hyped ChatGPT-5 turned out to be a dud — [incremental improvements wrapped in a routing architecture](#), nowhere near the breakthrough to AGI that Sam Altman had promised. Users are underwhelmed. As the *MIT Technology Review* [reports](#): “The much-hyped release makes several enhancements to the ChatGPT user experience. But it’s still far

² For example, [AI algorithms work very successfully to help landlords push the highest possible rents on tenants](#). Likewise, [the health insurance industry uses AI automation and predictive technologies](#) to systematically deny patients coverage for necessary medical care. AI also works for the military: defense company Anduril builds autonomous drones, virtual reality headsets, and other AI-powered technologies for the U.S. military. And private equity firms are hiring AI people to go through the companies they own and see how these should be restructured. In research, machine learning has revolutionized protein structure prediction and design, facilitating medicine development (see [Saplakoglu 2024](#) for a discussion).

short of AGI.” Instead, ChatGPT-5 makes [ridiculous errors](#), fails basic math and map problems, cannot count the number of fingers on one hand and ‘hallucinates’ as someone on a bad trip. The GPT-5 model [cannot correctly label Oregon on a map or correctly list the first twelve American presidents](#).

In fact, AI ‘hallucinations’ — the generation of fake citations or references³, non-existing historical events or people, plausible-sounding incorrect [legal advice](#) or [medical advice](#), or completely fabricated biologically-impossible images of animals or human beings — have become [more frequent in newer versions](#), supposedly designed for more complex tasks (see [here](#)). Worryingly, OpenAI’s internal tests show GPT-5 ‘hallucinates’ in circa one in 10 responses of the time on certain factual tasks, when it is connected to the internet. However, without web-browsing access, [GPT-5 is wrong in almost 1 in 2 responses](#), which should be troublesome. Even more worrisome, ‘hallucinations’ may also reflect [biases](#) buried within datasets. For instance, an LLM might ‘hallucinate’ crime statistics that align with racial or political biases simply because it has learned from biased data.

Of note here is that AI chatbots can be and are actively used to spread misinformation (see [here](#) and [here](#)). According to recent research, chatbots spread false claims when prompted with questions about controversial news topics 35% of the time — almost double the 18% rate of a year ago ([here](#)). “The worst performer was Inflection AI, which provided false claims to news prompts 57% of the time. The rate for Perplexity was 47% and for Meta and ChatGPT it was 40%” ([Woollacott 2025](#)). These bots can also be hacked. In 2016, [Microsoft launched](#) a chatbot named Tay on Twitter — then had to shut it down 16 hours later when Twitter users turned it into a racist, homophobic sexbot. In 2022, [Meta launched Galactica](#) trained on 48 million examples of scientific articles, websites, textbooks, lecture notes, and encyclopedias. Instead of helping scientists, Galactica mindlessly spat out biased and incorrect nonsense and, after massive criticism, was taken down after three days. And in July 2025, Grok, the AI-powered chatbot created by xAI, began pushing antisemitic tropes in its responses to some users’ queries, weeks after Musk said the chatbot needed to be rebuilt because some replies were “too politically correct” in Musk’s opinion. AI curates, orders, presents, and censors information, influencing interpretation and debate, while pushing dominant (average or preferred) viewpoints while suppressing alternatives, quietly removing inconvenient facts or making up convenient ones. The key issue is: Who controls the algorithms? Who sets the rules for the tech bros? It is evident that by making it easy to spread “realistic-looking” misinformation and biases and/or suppress critical evidence or argumentation, GenAI does and will have non-negligible societal costs and risks — which have to be counted when assessing its impacts.

³ A [recent study](#) from *Columbia Journalism Review's Tow Center for Digital Journalism* tested eight AI-driven search tools by providing direct excerpts from real news articles and asking the models to identify each article's original headline, publisher, publication date, and URL. They find that [the AI models incorrectly cited sources in more than 60 percent of these queries](#), raising serious concerns about their reliability in correctly attributing news content. “Perplexity provided incorrect information in 37 percent of the queries tested, whereas ChatGPTSearch incorrectly identified 67 percent (134 out of 200) of articles queried. Grok 3 demonstrated the highest error rate, at 94 percent” (source: [Benj Edwards 2025](#)).

Building larger LLMs is leading nowhere

The ChatGPT-5 debacle raises serious doubts and existential questions about whether the GenAI industry's core strategy of building ever-larger models on ever-larger data distributions has already hit a wall. Critics, including cognitive scientist Gary Marcus ([here](#) and [here](#)), have long argued that simply scaling up LLMs will not lead to AGI, and GPT-5's sorry stumbles do validate those concerns. It is becoming more widely understood that [LLMs are not constructed on proper and robust world models](#), but instead are built to autocomplete, based on sophisticated pattern-matching — which is why, for example, they still cannot even play chess reliably and continue to make mind-boggling errors with startling regularity. This is how [Noam Chomsky, Ian Roberts and Jeffrey Watumull \(2023\)](#) explain the issue:

“.... we know from the science of linguistics and the philosophy of knowledge that [LLMs] differ profoundly from how humans reason and use language. These differences place significant limitations on what these programs can do, encoding them with ineradicable defects. [...] The human mind is not, like ChatGPT and its ilk, a lumbering statistical engine for pattern matching, gorging on hundreds of terabytes of data and extrapolating the most likely conversational response or most probable answer to a scientific question. On the contrary, the human mind is a surprisingly efficient and even elegant system that operates with small amounts of information; it seeks not to infer brute correlations among data points but to create explanations. [...] Their deepest flaw is the absence of the most critical capacity of any intelligence: to say not only what is the case, what was the case and what will be the case — that's description and prediction — but also what is not the case and what could and could not be the case. Those are the ingredients of explanation, the mark of true intelligence.”

ChatGPT's problems are universal: other AI bots including xAI's Grok, DeepSeek, Google's Gemini and Anthropic's Claude are not faring any better. Failure is inherent in the way all LLMs are built: these supposedly “self-learning” [models do not induce abstract models](#) of the world that they can reliably engage with, even when those models are given to them (as with the rules of chess). A [June 2025 paper by researchers at Apple](#), titled “*The Illusion of Thinking*”, demonstrates that the latest, biggest LLMs, by their design, are incapable of reasoning, “over-think” easy problems, waste computational power on “wrong answers” (after finding the “correct” one), and completely collapse when the complexity of the issue under consideration increases and/or they have to generalize outside the space of their training data (Shojaee *et al.* 2025). To illustrate the point: the Apple researchers find that the largest billion-dollar generative AI system [cannot solve the “Tower of Hanoi” puzzle](#), even when it is given the [solution algorithm](#) and all it has to do is follow the steps.

Apple's conclusions are confirmed by the [paper](#) “*Is Chain-of-Thought Reasoning of LLMs a Mirage? A Data Distribution Lens*”, written by Arizona State University researchers Chengshuai Zhao, Zhen Tan, Pingchuan Ma, Dawei Li, Bohan Jiang, Yancheng Wang, Yingzhen Yang, and Huan Liu (2025). Their results demonstrate that “LLMs are not principled reasoners, but rather sophisticated simulators of reasoning-like text.” Ergo: Current LLMs are incapable of reasoning and creating new explanatory knowledge. Underscoring this conclusion, a [recent paper by scientists at MIT and Harvard](#), titled “*What Has a Foundation Model Found? Using Inductive Bias to Probe for*

World Models”, shows that even when trained on *all of physics*, LLMs fail to uncover even the *existing* generalized and universal physical principles underlying their training data. Specifically, Keyon Vafa, Peter Chang, Ashesh Rambachan and Sendhil Mullainathan (2025) note that LLMs that follow a “Kepler-esque” approach, can successfully predict the next position in a planet’s orbit, but fail to find the underlying explanation of Newton’s Law of Gravity (see [here](#)). Instead, they resort to fitting made-up rules, that allow them to successfully predict the planet’s next orbital position, but [these models fail](#) to find the force vector at the heart of Newton’s insight. The MIT-Harvard paper is explained in [this video](#). LLMs cannot and do not infer physical laws from their training data. Remarkably, they cannot even identify the relevant information from the internet. Instead, they make it up.

For anyone hoping that further scaling — building ever larger LLMs — would lead to better outcome, this is very bad news (see also [Che 2025](#)). Scaling will not lead to AGI (see [Marcus 1998](#) and [Marcus 2022](#)). By their design, LLMs remain stochastic parrots, repeating a lot but understanding nothing (Bender *et al.* 2021): they predict words based on probability, learning patterns from vast datasets scraped from the internet, books, and other sources (in gross violation of the rights of scientists, authors, photographers and artists), generating text or images based on patterns, not comprehension (and certainly not based on any intelligence). In other words, lacking understanding and an underlying mental model of the world, LLMs ‘hallucinate’ always and everywhere — even when they are seemingly generating outputs that make sense. That is also why many users complain that LLMs do not deliver adequate responses in their specific areas of expertise, although to the same users they seem to work in other domains in which these users are not experts.

The term ‘hallucination’ is already wrong, because it anthropomorphizes a phenomenon by attributing intent and consciousness to an algorithm, which it pertinently does not have. We are imagining a mind behind the output of AI — an inclination cleverly exploited by the AI industry to sell customers on AI. As [Michael Hitzik](#) explains in the *Los Angeles Times*:

“Indeed, as long ago as the 1960s, that phenomenon was noticed by Joseph Weizenbaum, the designer of the pioneering chatbot ELIZA, which replicated the responses of a psychotherapist so convincingly that even test subjects who knew they were conversing with a machine thought it displayed emotions and empathy. “What I had not realized,” [Weizenbaum wrote in 1976](#), “is that extremely short exposures to a relatively simple computer program could induce powerful delusional thinking in quite normal people.” Weizenbaum warned that the “reckless anthropomorphization of the computer” — that is, treating it as some sort of thinking companion — produced a “simpleminded view of intelligence.””

This is the “Eliza Effect” ([here](#)) and it is dangerous, especially since the people most drawn to AI tend to be those who understand the technology the least (as [research shows](#)).

AI models may generate incorrect output (read Marcus [here](#) and [here](#)), because they are trained on a certain dataset (however large). Questions asked outside of the data distribution on which the AI

model was trained, often leads to the bot responding with incorrect output. This happens because these models are prompted to respond to queries with the most statistically likely answer (given its training); hence, [the AI bot is incentivized to guess \(and give an incorrect response\) rather than admit it does not know something](#). This problem is recognized by researchers from OpenAI in a [recent paper](#). Guessing is rewarded — because, who knows, it *might* be right. The error is incorrigible. Accordingly, it might well be prudent to think of “Artificial Information” rather than “Artificial Intelligence” when using the acronym AI.

In sum, LLMs are super-expensive energy- and water-guzzling statistical pattern matchers, which are burning through billions of U.S. dollars to train the algorithms on particular data distributions, but which tend to generate out-of-distribution failures or break downs (because they do not have proper world models). Silicon Valley's endless appetite for funding was [put in an entirely new perspective earlier this year](#) when Chinese AI company DeepSeek demonstrated that its AI chatbot, DeepSeek V3, which was developed at a cost of less than \$6 million, could trade blows with the latest large language models made by OpenAI, Meta, and Google. The arrival of DeepSeek triggered [a more than \\$1 trillion sell-off](#) of tech stocks in January 2025. But the American AI industry recovered from this body-blow, by stepping up the hype and hyperbole around AGI and by begging for ever more cash. By refusing to admit defeat the industry took a leaf straight out of “*The Art of the Deal*”.

And adding insult to injury, the AI bots are founded on outright piracy. On September 5, 2025, [Anthropic agreed to pay](#) \$1.5 billion to settle claims that it downloaded pirated books to train its AI systems. The Anthropic settlement sends a chilling signal to all AI companies that downloading illegal copies of books, movies and images to train AI comes with a heavy cost. This will, for sure, lead to further settlements, given the potential non-negligible liability AI companies risk when they make use of books, movies and images from illegal channels (see [The Authors Guild 2025](#)). In February 2025, [Thomson Reuters won a copyright infringement lawsuit](#) against an unnamed AI company accused of using its proprietary content for AI model training without authorization. And Disney and Warner Brothers are suing AI companies for using images of copy-righted characters (think: Bugs Bunny and Superman).

On August 19, 2025, Eric Schmidt (yes, the same person who only four months earlier predicted that we would have AGI “within three to five years”) coauthored a [New York Times opinion article](#), arguing that “it is uncertain how soon artificial general intelligence can be achieved.” Schmidt now urges Silicon Valley to stop fixating on AGI, warning that the obsession distracts from building useful technology. This shows that Schmidt is capable of understanding (reality), unlike the LLMs.

95% of generative AI pilot projects in companies are failing

Corporations had rushed to announce AI investments or claim AI capabilities for their products in the hope of turbocharging their share prices. Then came the news that the AI tools are not doing what they are supposed to do and that people are realizing it (see [Ed Zitron](#)). An August 2025 [report](#) titled *The GenAI Divide: State of AI in Business 2025*, published by MIT's [NANDA](#) initiative,

concludes that 95% of generative AI pilot projects in companies are failing to raise revenue growth. (Intriguingly, MIT NANDA has now replaced the link to the report with a Google Form to request access, but — as [Ed Zitron](#) points out — [you can find the full PDF here](#)). As reported by [Fortune](#), “generic tools like ChatGPT [...] stall in enterprise use since they don’t learn from or adapt to workflows”. Quite.

Not mincing their words, the MIT NANDA authors Challapally, Pease, Raskar and Chari (2025) debunk “Five Myths About GenAI in the Enterprise”:

1. **AI will replace most jobs in the next few years** → Research found limited layoffs from GenAI, and only in industries that are already affected significantly by AI. There is no consensus among executives as to hiring levels over the next 3-5 years.
 2. **Generative AI is transforming business** → Adoption is high, but transformation is rare. Only 5% of enterprises have AI tools integrated in workflows at scale and 7 of 9 sectors show no real structural change.
 3. **Enterprises are slow in adopting new technology** → Enterprises are extremely eager to adopt AI and 90% have seriously explored buying an AI solution.
 4. **The biggest thing holding back AI is model quality, legal, data, risk** → What is really holding it back is that most AI tools don't learn and don't integrate well into workflows.
 5. **The best enterprises are building their own tools** → Internal builds fail twice as often.
- (Source: Challapally, Pease, Raskar and Chari 2025).

Indeed, firms are backpedaling after cutting hundreds of jobs and replacing these by AI. For instance, Swedish “*Buy Burritos Now, Pay Later*” Klarna bragged in March 2024 that its AI assistant was doing the work of (laid-off) 700 workers, only to rehire them (sadly, as gig workers) in the summer of 2025 (see [here](#)). Duolingo CEO Luis von Ahn [walked back on his earlier promises to replace all contract workers](#) with AI following a strong backlash from his customers. Other examples include IBM, forced to reemploy staff after laying off about 8,000 workers to implement automation (see this [The Wall Street Journal](#) article), and Australia’s Commonwealth Bank which is rehiring fired workers after its chatbot is found terrible at its job (see Marcus, [here](#)). Recent U.S. Census Bureau data by firm size show that AI adoption has been declining among companies with more than 250 employees (see **Figure 1**).

Tesla’s over-promised, under-delivered robo-taxi’s still need co-drivers (“a guy in the passenger seat”) and Tesla’s self-driving features are linked to numerous accidents ([here](#)), prompting safety concerns among both EV-AI enthusiasts and critics; in fact, it is impossible to call a car “smart” when its autopilot has been responsible for a series of shocking incidents – from drivers trapped in burning vehicles to dramatic stops on the highway (see [Sönke Iwersen and Michael Verfürden](#) (2025) in *The Guardian*). [The National Highway Traffic Safety Administration](#) opened an investigation into possible defects in Tesla doors (due to low battery voltage) that have reportedly left parents with children trapped in the back seat and forced to break windows to get them out.⁴


⁴ Robotaxi-firm Waymo which uses [LIDAR](#) and prioritizes safety and market-by-market mapping before expanding its services, [appears to work significantly better](#) than Tesla’s robotaxis.

As the [BBC reports](#), there is now a cottage industry of writers and coders who specialize in fixing AI's mistakes. What is more, the AI magic is built on the backs of thousands of (mostly underpaid) tech workers in back-offices who manually impose “guardrails” on AI output, rating and removing “weird” AI responses. Tellingly, as reported in *The Guardian* ([Bansal 2025](#)), “most [of these] workers said they avoid using LLMs or use extensions to block AI summaries because they now know how it’s built. Many also discourage their family and friends from using it, for the same reason.”

FIGURE 1

AI adoption rates starting to decline for larger firms

APOLLO



Source: Torsten Sløk (2025). <https://www.apolloacademy.com/ai-adoption-rate-trending-down-for-large-companies/> Note: The US Census Bureau conducts a biweekly survey of 1.2 million firms. Businesses are asked whether they have used AI tools such as machine learning, natural language processing, virtual agents or voice recognition to help produce goods or services in the past two weeks.

MIT economist [Daren Acemoglu \(2025\)](#) predicts rather modest productivity impacts of AI in the next 10 years and warns that some applications of AI may have negative social value. “We’re still going to have journalists, we’re still going to have financial analysts, we’re still going to have HR employees,” [Acemoglu says](#). “It’s going to impact a bunch of office jobs that are about data summary, visual matching, pattern recognition, etc. And those are essentially about 5% of the economy.” Similarly, using two large-scale AI adoption surveys (late 2023 and 2024) covering 11 exposed occupations (25,000 workers in 7,000 workplaces) in Denmark, Anders Humlum and

Emilie Vestergaard (2025) show, in a recent [NBER Working Paper](#), that the economic impacts of GenAI adoption are minimal: “AI chatbots have had no significant impact on earnings or recorded hours in any occupation, with confidence intervals ruling out effects larger than 1%. Modest productivity gains (average time savings of 3%), combined with weak wage pass-through, help explain these limited labor market effects.” These findings provide a much-needed reality check for the hyperbole that GenAI is coming for all of our jobs. Reality is not even close.

GenAI will not even make tech workers who do the coding redundant, contrary to the prediction by AI enthusiasts. Only in March 2025, Anthropic’s [Dario Amodei predicted](#) that in the next 3 to 6 months, AI would write 90% of all code, and in 12 months, nearly all computer code would be written by AI. [But it is not happening](#). Why not? [OpenAI researchers have found](#) (in early 2025) that advanced AI models (including GPT-4o and Anthropic's Claude 3.5 Sonnet) still are no match for human coders. The AI bots failed to grasp how widespread bugs are or to understand their context, leading to solutions that are incorrect or insufficiently comprehensive. Another [new study](#) from the nonprofit Model Evaluation and Threat Research (METR) finds that in practice, programmers, using early 2025-AI-tools, are actually *slower* when using AI assistance tools, spending 19 percent *more* time when using GenAI than when actively coding by themselves (see [here](#)). Programmers spent their time on reviewing AI outputs, prompting AI systems, and correcting AI-generated code.⁵

Using AI is also introducing new — unintended and significant — [security risks](#) into software development, which increases the need for tech workers to check software security. Untrained programmers writing and fixing code by describing what they want to an AI bot, may not only be introducing errors into their code, but also may be self-sabotaging by introducing [severe cybersecurity risks](#) to their code. Everyone using GenAI for coding should read [this important piece](#) by Gary Marcus and Nathan Hamiel (2025) on “*LLMs + Coding Agents = Security Nightmare*”. Gary Marcus ([here](#)) offers this as a summary:

“As a fantastic, souped-up version of autocomplete, AI coding tools somewhat increase productivity. They can help a coder learn a new API, or maybe even a new programming language. But current AI tools don’t replace an understanding of debugging, nor an understanding of system architecture, nor an understanding of what clients want. [...] The idea that coders (and more generally, software architects) are on their way out is absurd. AI will be a tool to help people write code, just as spell-checkers are a tool to help authors

⁵ However, according to Thomas Ferguson (see [French 2025](#)), AI may have led to declining employment of software programmers and some other fields. He points to a [graph](#) on the Federal Reserve Economic Data site that shows demand plunging virtually to zero in the last few years. However, the decline in software development job postings is not only due to AI, but (as Ferguson observes) is also caused by rising interest rates that have forced companies to adopt more conservative hiring practices. Outsourcing of software development to lower-cost regions such as India and Eastern Europe, and by a change in U.S. tax rules requiring software development costs to be amortized over just five years may also have contributed.

write articles and novels, but AI will not soon replace people who understand how to conceive of, write, and debug code.”

The U.S. economy at large is hallucinating


The disappointing rollout of ChatGPT-5 raises doubts about OpenAI’s ability to build and market consumer products that *users are willing to pay for*. That should be particularly concerning for financial investors, given that OpenAI, which has never made a profit so far ([here](#)), is reportedly in talks to sell shares, held by current and former employees, at a valuation of a whopping [\\$500 billion](#). “Five-hundred-unicorn” OpenAI would have a larger market capitalization than (say) ASML, Cisco Systems or Advanced Micro Devices (AMD), all established tech firms that are actually making profits (unlike OpenAI). The proposed secondary stock sale would not bring new capital into OpenAI but would instead allow firm insiders to cash out — and given how things are going, OpenAI employees may be well advised to sell their shares sooner than later.

But the point I want to make is not just about OpenAI: the American AI industry as a whole has been built on the premise that AGI is just around the corner. All that is needed is sufficient “compute”, *i.e.*, millions of Nvidia AI GPUs, enough data centers and sufficient cheap electricity to do the massive statistical pattern mapping needed to generate (a semblance of) “intelligence”. This, in turn, means that “scaling” (investing billions of U.S. dollars in chips and data centers) is the one-and-only way forward — and this is exactly what the tech firms, Silicon Valley venture capitalists and Wall Street financiers are good at: mobilizing and spending funds, this time for “scaling-up” generative AI and building data centers to support all the expected future demand for AI use.

During 2024 and 2025, Big Tech firms invested a staggering \$750 billion in data centers in cumulative terms and they plan to roll out a cumulative investment of \$3 trillion in data centers during 2026-2029 ([Thornhill 2025](#)). The so-called “Magnificent 7” (Alphabet, Apple, Amazon, Meta, Microsoft, Nvidia, and Tesla) spent more than \$100 billion on data centers in the second quarter of 2025; **Figure 2** gives the capital expenditures for four of the seven corporations.

The surge in corporate investment in “information processing equipment” is huge. According to Torsten Sløk, chief economist at Apollo Global Management, data center investments’ contribution to (sluggish) real U.S. GDP growth has been the same as consumer spending over the first half of 2025 (**Figure 3**). Financial investor Paul [Kedrosky](#) finds that capital expenditures on AI data centers (in 2025) have passed the peak of telecom spending during the dot-com bubble (of 1995-2000).


FIGURE 2



Christopher Mims (2025), <https://x.com/mims/status/1951256592642441239>

FIGURE 3

Contribution to GDP from consumer spending and data center investments




Source: Torsten Sløk (2025). <https://www.apolloacademy.com/similar-contribution-to-gdp-growth-from-consumer-spending-and-data-center-investments/>

Following the AI hype and hyperbole, tech stocks have gone through the roof, as shown in **Figure 4**. The S&P500 Index rose by circa 58% during 2023-2024, driven mostly by the growth of the share prices of the Magnificent Seven. The weighted-average share price of these seven

corporations increased by 156% during 2023-2024, while the other 493 firms experienced an average increase in their share prices of just 25%. America's stock market is largely AI-driven.

FIGURE 4
The S&P500 versus S&P7 versus S&P493
(January 2023 – December 2024)



Source: <https://get.ycharts.com/resources/blog/what-happened-to-the-magnificent-seven-stocks/>

Nvidia's shares rose by more than 280% over the past two years amid the exploding demand for its GPUs coming from the AI firms; as one of the most high-profile beneficiaries of the insatiable demand for GenAI, Nvidia now has a market capitalization of more than \$4 trillion, up by a factor 10 from around \$400 billion in mid-2022, which is the highest valuation ever recorded for a publicly traded company. Does this valuation make sense? Nvidia's price-earnings (P/E) ratio peaked at 234 in July 2023 and has since declined to 47.6 in September 2025 — which is still historically very high (see **Figure 5**). Nvidia is selling its GPUs to neocloud companies (such as CoreWeave, Lambda, and Nebius), which are funded by credit, from Goldman Sachs, JPMorgan, Blackstone and other Wall Street private equity firms, that is collateralized by the data centers filled with GPUs. [There is no real demand for AI compute](#), other than Microsoft, Meta, Amazon, OpenAI or Nvidia paying these companies to rent GPUs for “research.” In key cases, as explained by [Ed](#)

[Zitron](#), Nvidia offered the neocloud companies, which are loss making, to buy unsold cloud compute worth billions of U.S. dollars, effectively backstopping its clients — all in the expectation of an AI revolution that still has to arrive.

Likewise, the share price of Oracle Corp. (which is not included in the “Magnificent 7”) rose by more than 130% during mid-May and early September 2025 following the announcement of its \$300 billion cloud-computing infrastructure deal with OpenAI. Oracle’s P/E ratio shot up to almost 68, which means that financial investors are willing to pay almost \$68 for \$1 of Oracle’s future earnings. One obvious problem with this deal is that OpenAI doesn’t have \$300 billion; the company made a loss of \$15 billion during 2023-2025 and is projected to make a further cumulative loss of \$28 billion during 2026-2028 (see below). It is unclear and uncertain where OpenAI will get the money from. Ominously, Oracle needs to build the infrastructure for OpenAI before it can collect any revenue. If OpenAI cannot pay for the enormous computing capacity it agreed to buy from Oracle, which seems likely, Oracle will be left with the expensive AI infrastructure, for which it may not be able to find alternative customers, especially once the AI bubble fizzles out.


The P/E ratio of the other S&P493 firms (on September 5, 2025) is 15. The P/E ratios for the “Magnificent 7” are incredibly elevated (on September 5, 2025): Alphabet (25.4); Amazon (35.4); Apple (36.4); Meta (27.3); Microsoft (36.4); Nvidia (47.6); and Tesla (202.8). The elevated P/E ratios indicate that financial investors are willing to pay \$30 or \$40 for \$1 of earnings. Tesla’s P/E ratio looks like an especially poor bet in view of the fact that its EV sales are declining, as the company is getting pummeled by Chinese competitors including BYD, which now sells more cars globally than Tesla. Tesla is dropping prices for new vehicles in order to maintain customer demand, but to little effect; meanwhile, its robotaxis have failed to launch time and time again and the Cybertruck was a major flop. Perhaps Musk’s \$1 trillion pay package is more spectacle than reality, because the “ambitious” performance targets that he will have to meet, appear to be completely unattainable.

Tech stocks thus are considerably overvalued. Torsten Sløk, chief economist at Apollo Global Management, warned (in July 2025) that AI stocks are even more overvalued than dot-com stocks were in 1999. In a [blogpost](#), he illustrates how P/E ratios of Nvidia, Microsoft and eight other tech companies are higher than during the dot-com era (see **Figure 5**). We all remember how the dot-com bubble ended — and hence Sløk is right in sounding the alarm over the apparent market mania, driven by the “Magnificent 7” that are all heavily invested in the AI industry.

FIGURE 5

APOLLO

The AI bubble today is bigger than the IT bubble in the 1990s



Source: Torsten Sløk (2025), <https://www.apolloacademy.com/ai-bubble-today-is-bigger-than-the-it-bubble-in-the-1990s/>

Crucially, Big Tech does not buy these data centers and operate them itself, but instead the data centers are built by construction companies and then purchased by data center operators who lease them to (say) OpenAI, Meta or Amazon (see [here](#)). Wall Street private equity firms such as Blackstone and [KKR](#) are investing billions of U.S. dollars to buy up these data center operators, [using commercial mortgage-backed securities as source funding](#). Data center real estate is a new, hyped-up asset class that is beginning to dominate financial portfolios. Blackstone calls data centers one of its “highest conviction investments.” Wall Street loves the lease-contracts of data centers which offer long-term stable, predictable income, paid by AAA-rated clients like AWS, Microsoft and Google. Some Cassandra’s are warning of a potential oversupply of data centers, but given that “the future will be based on GenAI”, what could possibly go wrong?

AI-focused data centers guzzle enormous amounts of electricity, which [can be about 10 to 20 times](#) as much per server or rack as general cloud computing. These centers also need to be kept operational 99.999 percent of the day, which equates to about five minutes of downtime all year for maintenance or to switch out servers (source: [NYT](#)). The *U.S. Department of Energy* estimates that data centers (“hyper users”) already consumed 4.4% of U.S. electricity in 2023, and expects that to grow to as much as 12% by 2028 ([Paul Krugman](#)), driving up retail electricity prices for American households. Depending on how the electricity will be generated, the AI bubble may derail efforts to decarbonize the economy and lead to additional climate costs the world cannot afford. It must be noted that right now, the exponentially rising energy demands of the digital economy — sadly —

outpace the push to “decarbonize the grid”. Furthermore, some of the data centers are polluting their surroundings (see [Paul Krugman](#) and [Politico](#)), violating environmental laws. AI data centers which generate enormous heat, also consume scarce fresh water in two ways: for onsite server cooling and for offsite electricity generation. [One study](#) finds that generating a single 100-word email on ChatGPT uses around 1.5 liters of water ([Li et al. 2025](#)). In some U.S. regions where data centers are concentrated — and especially in those already facing shortages — [the strain on local water systems can be significant](#). *Bloomberg News* reports that about two-thirds of U.S. data centers built since 2022 [are in high water-stress areas](#).⁶

In a rare moment of frankness, OpenAI CEO Sam Altman has it right. “Are we in a phase where investors as a whole are overexcited about AI?” Altman said during [a dinner interview with reporters](#) in San Francisco in August. “My opinion is yes.” He also compared today’s AI investment frenzy to the dot-com bubble of the late 1990s. “Someone’s gonna get burned there, I think,” Altman said. “Someone is going to lose a phenomenal amount of money – we don’t know who ...”, but (going by what happened in earlier bubbles) it will most likely not be Altman himself.

The question therefore is: How long investors will continue to prop up sky-high valuations of the key firms in the GenAI race remains to be seen. Earnings of the AI industry continue to pale in comparison to the tens of billions of U.S. dollars that are spent on data center growth. According to an upbeat [S&P Global research note](#) published in June, 2025 the GenAI market is projected to generate \$85 billion in revenue 2029. However, Alphabet, Google, Amazon and Meta together will [spend nearly \\$400 billion](#) on capital expenditures in 2025 alone, or \$520 billion according to [Kedrosky](#). At the same time, the AI industry has a combined revenue that is little more than the revenue of the smart-watch industry (see [Ed Zitron 2025](#)).

America’s AI industry is non-transparent — OpenAI and Anthropic do not share numbers on realized revenues and costs but instead obfuscate their actual performance. Their capital expenditures go through the roof (**Figure 2**) and their revenues remain unclear. It is uncertain whether the massive bet on AI will eventually pay off. Technology writer and AI critic Ed Zitron has [compared the situation](#) to the subprime mortgage crisis of 2007-08. Zitron has collected and analyzed the relevant numbers ([here](#) and [here](#)) and the emerging picture is not good. For instance, OpenAI raised and spent around \$40 billion so far and intends to receive around \$22.5 billion by the end of year from SoftBank; [per Bloomberg](#), OpenAI expects to earn a revenue of \$12.7 billion in 2025 (compared to around \$4.3 billion in 2024). [OpenAI made a loss](#) of \$2 billion in 2023, \$5 billion in 2024, and \$8 billion in 2025 (projected) — and another projected \$28 billion loss during 2026-2028. Amodeli’s Anthropic raised \$20.7 billion so far, plus \$5 billion more this year; and its revenues in 2024 and 2025 are \$918 million and (a projected) \$2-4 billion, respectively. Anthropic is losing \$3 billion in 2025. [Musk’s xAI is currently spending](#) \$1 billion every month and expects to spend a staggering \$13 billion in 2025. xAI is only expecting revenues worth \$500 million in 2025, hoping these will climb to over \$2 billion in 2026. Nonetheless, Musk claims that xAI will turn a

⁶ The *Financial Times* published a useful report on pressures on Big Tech to act on data centers’ thirst for water; see [Temple-West \(2025\)](#).

profit in 2027, somehow. Mind you, even these dreadful numbers are most likely too optimistic, as [Zitron documents](#) in detail.⁷

On top of this, computing costs are rising (not declining), making life even more difficult for the industry. The newer AI models try to imitate “reasoning” by breaking a user's input into component parts, then run inference on each one of those parts — which means, in the case of Grok 4, using 603 tokens to generate *two words*. In effect, with a steady price per token and more token consumption per word, AI inference costs have grown about 10x over the last two years (see the analysis by Ewa Szyszka [here](#)). Cost of GenAI applications therefore go up. Nonetheless, AI firms expect to be profitable, somehow, by 2030, latest.

So, what if GenAI just is not profitable? This question is pertinent in view of the rapidly diminishing returns to the stratospheric capital expenditures on GenAI and data centers and the disappointing user-experience of 95% of firms that adopted AI. However, it is easier to fool people than to convince them that they have been fooled.⁸ Only being honest, one of the largest hedge funds in the world, Florida-based [Elliott](#), told clients that AI is overhyped and Nvidia is in a bubble, adding that many AI products are “never going to be cost-efficient, never going to actually work right, will take up too much energy, or will prove to be untrustworthy.” “There are few real uses,” it said, other than “summarizing notes of meetings, generating reports and helping with computer coding”. It added that it was “sceptical” that [Big Tech companies would keep buying](#) the chipmaker’s graphics processing units in such high volumes.

As argued by [Zitron](#) and others, LLMs do not scale profitably and their functionality is not significant enough to justify the costs of running them. OpenAI is estimated to have 700 million weekly active users — but how many of them are paying? Zitron’s data suggest that OpenAI has 20 million paying ChatGPT users and 5 million “business customers”. This is already a mediocre (3.5%) conversion rate (see Zitron [here](#)). And what is the churn rate? How many paid users are ending their subscription? Locking billions of U.S. dollars in into AI-focused data centers without a clear exit strategy for these investments in case the AI craze ends, only means that systemic risk in finance and the economy is building. With data-center investments driving U.S. economic growth, the American economy has become dependent on a handful of corporations, which have not yet managed to generate a profit on these data center investments.

Now that GenAI is hitting a wall, the U.S. economy may have reached a breaking point: soon it may become clear that one cannot live in a data center. As [Paul Kedrosky](#) puts it: “Regardless of what one thinks about the merits of AI or explosive datacenter expansion, the scale and pace of capital deployment into a rapidly depreciating technology is remarkable. These are not railroads — we aren’t building century-long infrastructure. AI datacenters are short-lived, asset-intensive

⁷ The [United Arab Emirates](#), which has committed to a 10-year, \$1.4 trillion investment framework in the United States in AI infrastructure, semiconductors, energy, and manufacturing, will hold some of the bag, once things go wrong.

⁸ The quip is attributed to Mark Twain who, however, did not write it; instead, Twain wrote that “How easy it is to make people believe a lie, and how hard it is to undo that work again!”

facilities riding declining-cost technology curves, requiring frequent hardware replacement to preserve margins.”

Kedrosky may well be too pessimistic, because parts of the data-center infrastructure may turn out to be useful even following an AI spending crash, but he is right in pointing out an essential difference between railroad infrastructure and digital infrastructure: railroads have a useful economic life of 40-50 years (and a corresponding straight-line annual depreciation rate of 2%-3%), whereas servers, networking equipment and storage devices have a useful lifetime of 3-5 years (and a corresponding annual depreciation rate of 20%-30%).

America’s high-stakes geopolitical bet gone wrong

The AI boom (bubble) developed with the support of both major political parties in the U.S. The vision of American firms pushing the AI frontier and reaching GenAI first is widely shared — in fact, there is a bipartisan consensus on how important it is that the U.S. should win the global AI race. When President Donald Trump announced “[Winning the AI Race: America’s AI Action Plan](#)” on July 2025, he said something surprisingly uncontroversial: “America is the country that started the AI race. And as president of the United States, I’m here today to declare that America is going to win it.”

Winning the AI race with China is regarded to be of existential importance to the U.S. economy and its geopolitical standing in the increasingly belligerent multi-polar world order, in which the U.S. no longer has the power, as before, to unilaterally use the weapons of economic, financial and technological coercion — sanctions, access to global payment systems, access to U.S. dollar loans, internet surveillance, supply chain interventions, and export controls — for its own strategic advantage and national security interests (see Farrell and Newman 2025). Today, other nations, most prominently China, have the means to use similar economic and financial weapons to counter U.S. power.

The real loss of U.S. control is illustrated by the fact that the Trump administration was forced to make concessions on semiconductor export controls in return for China’s easing restrictions on rare-earth minerals that were crippling the American auto industry. China effectively weaponized its near-monopoly in these minerals, adapting very quickly to the new world order. The lesson for the U.S. came through loud and clear (Farrell and Newman 2025): America’s industrial capability is critically dependent on a number of potential adversary nation-states, including China, as Secretary of State [Marco Rubio acknowledged](#) in June 2025. In this context, America’s lead in GenAI is considered to constitute a potential very powerful geopolitical lever: If America manages to get to AGI first, so the analysis goes, it can build up an overwhelming long-term advantage over especially China.

The Chinese appear to agree, but, crucially, they are taking a radically different approach to the issue. [Song-Chun Zhu, a top Chinese AI scientist](#) and director of the *Beijing Institute for General Artificial Intelligence*, insists that Silicon Valley’s current quest for AGI is built on sand. “A sign of

true intelligence, he argues, is the ability to reason towards a goal with minimal inputs – what he calls a “small data, big task” approach, compared with the “big data, small task” approach employed by LLMs like ChatGPT. AGI, Zhu’s team has recently said, is characterized by qualities such as resourcefulness in novel situations, social and physical intuition, and an understanding of cause and effect. Large language models, Zhu believes, will never achieve this” (source: [Che 2025](#)).

For [Rubio](#), “winning the AI race is non-negotiable. America must continue to be the dominant force in artificial intelligence to promote prosperity and protect our economic and national security.” To do so, according to this view, the U.S. should use its chokehold on advanced semiconductors to keep the control over the development of AI, slow down China’s build-up of AGI, and thereby secure America’s long-term dominance. Reaching AGI first is thus considered instrumental in shaping the world in accordance with the view of Silicon Valley, Wall Street and U.S. national security policymakers.

This bipartisan America-centric (or “America First”) view is based on four assumptions (per Henry Farrell [here](#)): “that competition between the US and China is everything, that AGI is right around the corner, that whoever gets to AGI first is likely to win, and that America’s big advantage is its chokehold over the chips that you need to train powerful AI.” This is not just the belief of the Trump administration, but it also underpinned Biden’s approach to China. The idea, influenced by billionaire end-of-times Silicon Valley prophets, was that the nation that would reach AGI first would enjoy structural economic *and military* advantages that would help to secure its long-term geopolitical dominance ([Farrell again](#)). Following the cold douche of August 2025, that strategy is in shreds.

The administration appears to be in denial. On September 4th, President Trump gathered a who’s who of the [American tech elite for dinner](#) at the White House: Meta’s Mark Zuckerberg; Microsoft’s Bill Gates; Apple CEO Tim Cook; Google’s Sergey Brin and Sundar Pichai; and OpenAI’s Sam Altman. (Could not make it: Elon Musk). All the dinner guests put Uriah Heep in the shade by a sycophantic display of fealty to the President. Meta’s Zuckerberg and Apple’s Tim Cook each pledged to invest “at least \$600 billion” in the US over the next few years in order to build their data centers and infrastructure to power the next wave of innovation. “That’s a lot,” Trump responded upon hearing the amounts.

He’s right. Silicon Valley, Wall Street and the Trump administration are doubling down on the “AGI First” strategy. But more astute observers did note its collapse — and highlight the cost of denialism. Prominently, [Eric Schmidt and Selina Xu](#) worry, in the *New York Times* of August 19, 2025, that “Silicon Valley has grown so enamored with accomplishing this goal [of AGI] that it’s alienating the general public and, worse, bypassing crucial opportunities to use the technology that already exists. In being solely fixated on this objective, our nation risks falling behind China, which is far less concerned with creating A.I. powerful enough to surpass humans and much more focused on using the technology we have now.”

This is the point indeed. Unlike the U.S., China has not bet the farm on GenAI or AGI, but instead has been building its own alternative cluster of mutually reinforcing high-tech industries centered

on the (renewable) energy economy and AI (see [Che 2025](#)). As argued and documented by [Tim Sahay and Kate MacKenzie \(2025\)](#), China is leading in renewable and electrification technologies, ranging from electric vehicles, to batteries to solar panels, industrial robots, and software — while it is also controls the supply chains of many critical raw materials. China recently banned its tech companies from buying Nvidia’s AI chips, as Beijing steps up its efforts to boost semiconductor independence. Other (developing) countries exit the orbit of the U.S. and turn to the Chinese for trade and foreign investment. The U.S. is falling behind in exactly these areas, although it still maintains its hold on the dollar-based global financial system, global intellectual property and on global technology platforms (the internet and much of social media). Doubling down on AGI will help neither the U.S. economy nor its geopolitical standing, however — the high-stakes on the “big data, small tasks” AGI bet has already gone wrong.

Perhaps the plight of the U.S. economy is captured best by OpenAI’s Sam Altman who fantasizes about putting his data centers in space: “Like, maybe we build a big Dyson sphere around the solar system and say, “Hey, it actually makes no sense to put these on Earth.”” For as long as such [‘hallucinations’ on using solar-collecting satellites to harvest \(unlimited\) star power](#) continue to convince gullible financial investors, the government and users of the “magic” of AI and the AI industry, the U.S. economy is surely doomed.

References

- Acemoglu, D. 2025. “The simple macroeconomics of AI.” *Economic Policy* 40(121): 13–58, <https://doi.org/10.1093/epolic/eiae042>
- Al-Sibai, N. 2025. ‘OpenAI Researchers Find That Even the Best AI Is "Unable To Solve the Majority" of Coding Problems.’ *Futurism*, February 23. <https://futurism.com/openai-researchers-coding-fail>
- Al-Sibai, N. 2025. ‘What Actually Happens When Programmers Use AI Is Hilarious, According to a New Study.’ *Futurism*, July 16. <https://futurism.com/ai-coding-programmers-reality>
- Al-Sibai, N. 2025. ‘Programmers Using AI Create Way More Glaring Security Issues, Data Shows.’ *Futurism*, September 8. <https://futurism.com/ai-coding-security-problems>
- Amaral, C. 2024. ‘AI lending will make finance deals even more unfair for women – here’s how this can be avoided.’ *The Conversation*, August 30. <https://doi.org/10.64628/AB.y6vtnfqw4>
- Authors Guild. 2025. ‘What Authors Need to Know About the \$1.5 Billion Anthropic Settlement.’ September 5. <https://authorsguild.org/news/what-authors-need-to-know-about-the-anthropic-settlement/>
- Bansal, V. 2025. ‘How thousands of ‘overworked, underpaid’ humans train Google’s AI to seem smart.’ *The Guardian*, September 11. <https://www.theguardian.com/technology/2025/sep/11/google-gemini-ai-training-humans>
- Becker, J. N. Rush, E. Barnes and D. Rein. 2025. ‘Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.’ July 10. METR Model Evaluation & Threat Research. <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>
- Bender, E.M., T. Gebru, A. McMillan-Major and S. Shmitchell. 2021. ‘On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?’  [FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency](https://doi.org/10.1145/3442188.3445922) pp. 610 – 623 <https://doi.org/10.1145/3442188.3445922>
- Bubeck, S. *et al.* 2023. ‘Sparks of Artificial General Intelligence: Early experiments with GPT-4.’ Research Paper. <https://doi.org/10.48550/arXiv.2303.12712>
- Challapally, A. *et al.* 2025. *The GenAI Divide. State of AI in Business 2025*. July. Project NANDA. Cambridge, Mass.: MIT.

- Che, C. 2025. "I have to do it": Why one of the world's most brilliant AI scientists left the US for China.' *The Guardian*, September 16. <https://www.theguardian.com/news/ng-interactive/2025/sep/16/song-chun-zhu-why-one-of-the-worlds-most-brilliant-ai-scientists-left-the-us-for-china>
- Chomsky, N. I. Roberts and J. Watumull. 2023. 'Noam Chomsky: The False Promise of ChatGPT.' *The New York Times*, March 8. <https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html>
- Dizikes, P. 2025. 'A Nobel laureate on the economics of artificial intelligence.' *MIT Technology Review*, February 25. <https://www.technologyreview.com/2025/02/25/1111207/a-nobel-laureate-on-the-economics-of-artificial-intelligence/>
- Estrada, E. 2025. 'MIT report: 95% of generative AI pilots at companies are failing.' *Fortune*, August 18. <https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/?ref=wheresyoured.at>
- Farrell, H. 2025. 'The twilight of tech unilateralism.' *Programmable Mutter*, Substack. August 19. <https://www.programmablemutter.com/p/the-twilight-of-tech-unilateralism>
- Farrell, M. 2025. 'Wall St. Is All In on A.I. Data Centers. But Are They the Next Bubble?' *The New York Times*, June 2. <https://www.nytimes.com/2025/06/02/business/ai-data-centers-private-equity.html>
- Farrel, H. and A. Newman. 2025. 'The Weaponized World Economy.' *Foreign Affairs* September/October 2025.
- Fletcher, L. and C. Mourselas. 2025. 'Elliott says Nvidia is in a 'bubble' and AI is 'overhyped'.' *The Financial Times*, August 2. <https://www.ft.com/content/24a12be1-a973-4efe-ab4f-b981aee0cd0b>
- French, N. 2025. 'Tech Capital Is Dominating American Politics. An Interview with Thomas Ferguson.' *Jacobin*, September 17. <https://jacobin.com/2025/09/tech-capital-american-politics-ferguson>
- Hiltzik, M. 2025. 'Say farewell to the AI bubble, and get ready for the crash.' *Los Angeles Times*, August 20. <https://www.latimes.com/business/story/2025-08-20/say-farewell-to-the-ai-bubble-and-get-ready-for-the-crash>
- Huckins, G. 2025. 'GPT-5 is here. Now what?' *MIT Technology Review*, August 7. <https://www.technologyreview.com/2025/08/07/1121308/gpt-5-is-here-now-what/>
- Humlum, A. and E. Vestergaard. 2025. 'Large Language Models, Small Labor Market Effects.' NBER Working Paper 33777. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w33777>

- Iwersen, S. and M. Verfürden. 2025. 'The vehicle suddenly accelerated with our baby in it': the terrifying truth about why Tesla's cars keep crashing.' *The Guardian*, July 5.
<https://www.theguardian.com/technology/2025/jul/05/the-vehicle-suddenly-accelerated-with-our-baby-in-it-the-terrifying-truth-about-why-teslas-cars-keep-crashing>
- Jaźwińska, K. and A. Chandrasekar. 2025. 'AI Search Has A Citation Problem.' *Columbia Journalism Review*, March 6. https://www.cjr.org/tow_center/we-compared-eight-ai-search-engines-theyre-all-bad-at-citing-news.php
- Li, P., J. Yang, M.A. Islam and S. Ren 2025. 'Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models.' Last revised version.
<https://arxiv.org/abs/2304.03271>
- Lin, B. 2025. 'IBM CEO Says AI Has Replaced Hundreds of Workers but Created New Programming, Sales Jobs.' *The Wall Street Journal*. May 6. <https://www.wsj.com/articles/ibm-ceo-says-ai-has-replaced-hundreds-of-workers-but-created-new-programming-sales-jobs-54ea6b58>
- Kedrosky, P. 2025. 'Honey, AI Capex is Eating the Economy.' July 18.
<https://paulkedrosky.com/honey-ai-capex-ate-the-economy/?ref=wheresyoured.at>
- Krugman, P. 2025. 'AI Is Power-Hungry.' Substack, August 18.
<https://paulkrugman.substack.com/p/ai-is-power-hungry>
- Krugman, P. 2025. 'What Happens If AI Hits An Energy Wall?' Substack, August 19.
<https://paulkrugman.substack.com/p/what-happens-if-ai-hits-an-energy>
- Mauran, C. 2025. 'OpenAI says GPT-5 hallucinates less — what does the data say?' Mashable.com, August 7. <https://mashable.com/article/openai-gpt-5-hallucinates-less-system-card-data>
- Marcus, G. 1998. 'Rethinking Eliminative Connectionism.' *Cognitive Psychology* 37 (3): 243-282.
<https://www.sciencedirect.com/science/article/pii/S0010028598906946>
<https://doi.org/10.1006/cogp.1998.0694>
- Marcus, G. 2022. 'Deep Learning Is Hitting a Wall.' *Nautilus*, March 10.
- Marcus, G. 2025. 'A knockout blow for LLMs?' *Marcus on AI*, Substack, June 8.
- Marcus, G. 2025. 'AI Coding Fantasy meets Pac-Man.' *Marcus on AI*, Substack, March 12.
https://garymarcus.substack.com/p/ai-coding-fantasy-meets-pac-man?utm_source=publication-search

- Marcus, G. 2025. ‘Generative AI’s crippling and widespread failure to induce robust models of the world.’ *Marcus on AI*, Substack, June 25. <https://garymarcus.substack.com/p/generative-ais-crippling-and-widespread?r=8tdk6>
- Marcus, G. 2025. ‘GPT-5: Overdue, overhyped and underwhelming. And that’s not the worst of it.’ *Marcus on AI*, Substack, August 10. <https://garymarcus.substack.com/p/gpt-5-overdue-overhyped-and-underwhelming>
- Marcus, G. 2025. ‘LLMs are not like you and me—and never will be.’ *Marcus on AI*, Substack, August 13. <https://garymarcus.substack.com/p/llms-are-not-like-you-and-meand-never>
- Marcus, G. 2025. ‘AI, layoffs, productivity and The Klarna Effect.’ *Marcus on AI*, Substack, August 23. <https://garymarcus.substack.com/p/ai-layoffs-productivity-and-the-klarna>
- Marcus, G. and N. Hamiel. 2025. ‘LLMs + Coding Agents = Security Nightmare.’ *Marcus on AI*, Substack, August 17. <https://garymarcus.substack.com/p/llms-coding-agents-security-nightmare>
- Metz, C. and Karen Weise. 2025. ‘A.I. is getting more powerful, but its hallucinations are getting worse.’ *The New York Times*, May 5. <https://www.nytimes.com/2025/05/05/technology/ai-hallucinations-chatgpt-google.html>
- Miserendino, S. *et al.* 2025. ‘SWE-Lancer: Can Frontier LLMs Earn \$1 Million from Real-World Freelance Software Engineering.’ May 29. <https://doi.org/10.48550/arXiv.2502.12115v4>
- Mitchell, H. 2025. ‘The Less You Know About AI, the More You Are Likely to Use It.’ *The Wall Street Journal*, September 2. <https://www.wsj.com/tech/ai/ai-adoption-study-7219d0a1>
- Morrow, A. 2025. ‘The latest ChatGPT is supposed to be ‘PhD level’ smart. It can’t even label a map.’ CNN Business, August 14. <https://edition.cnn.com/2025/08/14/business/chatgpt-rollout-problems>
- Romero, A. 2025. ‘Harvard and MIT Study: AI Models Are Not Ready to Make Scientific Discoveries.’ *The Algorithmic Bridge*, Substack, July 15. <https://www.thealgorithmicbridge.com/p/harvard-and-mit-study-ai-models-are?ref=wheresyoured.at>
- Sahau, T. and K. MacKenzie. 2025. ‘BRICS in 2025.’ *Phenomenal World*, July 5. <https://www.phenomenalworld.org/analysis/brics-in-2025/>
- Schmidt, E. 2025. ‘The AI revolution is underhyped.’ TED Talk, April. https://www.ted.com/talks/eric_schmidt_the_ai_revolution_is_underhyped
- Schmidt, E. and S. Xu. 2025. ‘Silicon Valley is drifting out of touch with the rest of America.’ *The New York Times*, August 19. <https://www.nytimes.com/2025/08/19/opinion/artificial-general-intelligence-superintelligence.html>

S&P Global. 2025. 'Generative AI market revenue projected to grow at a 40% CAGR from 2024–2029.' S&P Global, June 3. <https://www.spglobal.com/market-intelligence/en/news-insights/research/generative-ai-market-revenue-projected-to-grow-at-a-40-cagr-from-2024-2029>

Saplakoglu, Y. 2024. 'How AI Revolutionized Protein Science, but Didn't End It.' *Quanta Magazine*, June 26. <https://www.quantamagazine.org/how-ai-revolutionized-protein-science-but-didnt-end-it-20240626/>

Shojaee, P. *et al.* 2025. 'The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity.' Apple Co. June. <https://machinelearning.apple.com/research/illusion-of-thinking>

Sløk, T. 2025. 'AI Bubble Today Is Bigger Than the IT Bubble in the 1990s.' *The Daily Spark*, Apollo Global Management. July 16. <https://www.apolloacademy.com/ai-bubble-today-is-bigger-than-the-it-bubble-in-the-1990s/>

Sløk, T. 2025. 'Similar Contribution to GDP Growth From Consumer Spending and Data Center Investments.' *The Daily Spark*, Apollo Global Management. August 18. <https://www.apolloacademy.com/similar-contribution-to-gdp-growth-from-consumer-spending-and-data-center-investments/>

Sløk, T. 2025. 'AI Adoption Rate Trending Down for Large Companies.' *The Daily Spark*, Apollo Global Management. September 7. <https://www.apolloacademy.com/ai-adoption-rate-trending-down-for-large-companies/>

Szyszk, E. 2025. 'Future AI bills of \$100k/yr per dev.' *Kilo Code Blog*, Substack. August 8. <https://blog.kilocode.ai/p/future-ai-spend-100k-per-dev?ref=wheresyoured.at>

Tangermann, V. 2025. 'Companies Are Discovering a Grim Problem With "Vibe Coding"'. *Futurism*, May 31. <https://futurism.com/problem-vibe-coding>

Tangermann, V. 2025. 'OpenAI Realizes It Made a Terrible Mistake'. *Futurism*, September 14. <https://futurism.com/openai-mistake-hallucinations>

Tarnoff, B. 2023. 'Weizenbaum's nightmares: how the inventor of the first chatbot turned against AI.' *The Guardian*, July 25. <https://www.theguardian.com/technology/2023/jul/25/joseph-weizenbaum-inventor-eliza-chatbot-turned-against-artificial-intelligence-ai>

Tauman Kalai, A., O. Nachum, S.S. Vempala and E. Zhang. 2025. 'Why Language Models Hallucinate.' September 4, <https://arxiv.org/abs/2509.04664>

Temple-West, P. 2025. 'Big Tech under pressure to act on data centres' thirst for water.' *The Financial Times*, March 21. <https://www.ft.com/content/65fff689-bd47-4c15-bdb8-083e5ccd84dc>

- Thornhill, J. 2025. 'Brace for a crash before the golden age of AI.' *The Financial Times*, August 21. <https://www.ft.com/content/a76f238d-5543-4c01-9419-52aaf352dc23>
- Vafa, K. *et al.* 2025. 'What Has a Foundation Model Found? Using Inductive Bias to Probe for World Models.' <https://doi.org/10.48550/arXiv.2507.06952>
- Weizenbaum, J. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. The Johns Hopkins University Press.
- White House. 2025. 'White House Unveils America's AI Action Plan.' July 23. <https://www.whitehouse.gov/articles/2025/07/white-house-unveils-americas-ai-action-plan/>
- Wilkins, J. 2025. 'Exactly Six Months Ago, the CEO of Anthropic Said That in Six Months AI Would Be Writing 90 Percent of Code Not even close.'" *Futurism*, September 10. <https://futurism.com/six-months-anthropic-coding>
- Wittenberg, A. 2025. 'How come I can't breathe?': Musk's data company draws a backlash in Memphis.' *Politico*, June 5. <https://www.politico.com/news/2025/05/06/elon-musk-xai-memphis-gas-turbines-air-pollution-permits-00317582>
- Woollacott, E. 2025. 'AI Chatbots Are Feeding You More False Information Than Ever.' *Forbes*, September 5. <https://www.forbes.com/sites/emmawoollacott/2025/09/05/ai-chatbots-are-feeding-you-more-false-information-than-ever/>
- Zhao, C. *et al.* 2025. 'Is Chain-of-Thought Reasoning of LLMs a Mirage? A Data Distribution Lens.' <https://doi.org/10.48550/arXiv.2508.01191>
- Zitron, E. 2024. 'The Subprime AI CRisis.' *Where's Your Ed At?* Substack. September 16. <https://www.wheresyoured.at/subprimeai/>
- Zitron, E. 2025. 'How Much Money Do OpenAI And Anthropic Actually Make?' *Where's Your Ed At?* Substack. August 1. <https://www.wheresyoured.at/howmuchmoney/>
- Zitron, E. 2025. 'AI is a money trap.' *Where's Your Ed At?* Substack. August 6. <https://www.wheresyoured.at/ai-is-a-money-trap/>
- Zitron, E. 2025. 'How to argue with an AI booster.' *Where's Your Ed At?* Substack. August 25. <https://www.wheresyoured.at/how-to-argue-with-an-ai-booster/>
- Zitron, E. 2025. 'Is There Any Real Money In Renting Out AI GPUs?.' *Where's Your Ed At?* Substack. September 19. <https://www.wheresyoured.at/is-there-any-real-money-in-renting-out-ai-gpus/#there-is-no-real-market-for-ai-compute>

