

# The False Prophets of Silicon Valley

Artificial intelligence will transform the world — but not into the utopia its evangelists promise.

Kyle Hiebert

“In the future, everyone’s lives can be better than anyone’s life is now,” wrote OpenAI’s leader Sam Altman in a September [blog post](#). It’s a spectacular assertion from a 39-year-old CEO rumoured to be gaining a seven percent [equity stake](#) in a company valued at US\$157 billion. Altman insists artificial intelligence (AI) will soon outperform the mental abilities of most people. Afterward, he foresees climate change being solved, poverty erased and humanity colonizing space. All of this, Altman claims, is just around the corner — if only tech titans receive enough resources and leeway.

Such views are common in Silicon Valley. Google DeepMind CEO and co-founder Demis Hassabis [says](#) software capable of human-level reasoning — artificial general intelligence, or AGI — will emerge by 2030. DeepMind’s chief AGI scientist, Shane Legg, [pegs it at 2028](#). The head of Anthropic, Dario Amodei, sees humans lagging behind machines [as early as 2026](#). Elon Musk thinks this could happen by the [end of this year](#).

But these epic forecasts clash with the now grinding pace of AI breakthroughs. Meanwhile, evidence of the technology’s risks and intrinsic limitations is mounting. And the pitfalls on its path to profitability are multiplying. Increasingly, industry hype around AI delivering universal prosperity is being punctured by the [real-world liabilities](#) of its development.

Indeed, the AI fever dreams emanating from the walled garden of Silicon Valley probably aren’t realizable. Rather, they are avatars for the interests and values of the tech community’s ultra elite. They reflect a desire to re-engineer the human project.

In trying to manifest computer superintelligence, AI accelerationists seem to have overlooked a few key factors. Their bespoke visions for the future discount the vagaries

of consumer behaviour and state power, for example. Energy infrastructure is lacking. And companies aren't uniformly keen to disrupt their existing business models. Complicating matters further are trade wars, surging protectionism, finite global resources and knotted supply chains.

This is not to diminish the vital importance and vast promise of the technology. Machine-learning applications are already spurring much-needed innovation and productivity gains in nearly every sector of the global economy.

Researchers now have a powerful new tool to devise more effective vaccines and cancer treatments, thanks to the Nobel Prize-winning work of Hassabis and his colleague in using AI to solve the biological riddles of protein folding — a process that had long vexed medical science. Developing nations, too, stand to benefit enormously from AI's use to improve **population mapping**. These nations will also see a massive unlocking of human capital through the growth of AI-powered education and health information apps on smartphones.

New digital agriculture tech will be critical as well. Global food production must increase by an estimated

**60–70 percent** by mid-century to meet the demands of a forecasted population of 10 billion people. Embodying AI in robots can help address labour deficits that stem from aging workforces. Smarter weather models are already identifying the onset of ferocious storms much earlier than previous predictive systems, bolstering disaster relief. The delivery of public services by various levels of government is being enhanced through faster analysis of big data.

Even in armed conflict, there are some reasons to be optimistic. Ukraine has harnessed intelligent drones and algorithmic targeting software to defend itself against a menacing Russia. Military applications of AI no doubt represent a major arms control dilemma, given the possibility of humans becoming passive participants in war. But the deft adoption of **autonomous weapons systems** and AI-driven cyber defences by liberal democracies could also be key to deterring hostile autocracies in a more unstable world.

And AI will continue to improve; it must.

The internet is becoming saturated with AI content, hobbling the development of new foundational models. That's because of the shrinking amount of novel high-quality data that tech firms can scrape from the Web in the form of human-generated text, images and videos.

That said, the technology's advancements likely won't keep pace with the seamless upward trajectory envisioned by its most ardent cheerleaders. Public **trust in AI systems** is already trending downward globally. Generative AI models remain especially addled by profound inaccuracies and bizarre hallucinations. And such flaws appear entrenched for at least two reasons — data scarcity and economics. First, the internet is becoming saturated with AI content, hobbling the development of new foundational models. That's because of the shrinking amount of novel high-quality data that tech firms can scrape from the Web in the form of human-generated text, images and videos. AI systems become unstable after cannibalizing other machine-generated inputs. One **study** suggests that tech companies may exhaust the library of publicly available human text online sometime between 2026 and 2032.

It's always possible developers will innovate their way around this roadblock. But it's a long shot. Even the world's leading AI scientists admit they can't comprehend the inner workings of their creations. Anthropic has made the most progress, releasing a **research paper** this past May that identified how millions of concepts are represented within a version of Claude, their flagship large language model (LLM). However, that still leaves a lot to be desired. The concept-building algorithms of rival OpenAI's GPT-4, for example, involve an estimated 1.8 trillion parameters. In their paper, Anthropic's researchers confess that gaining a full understanding of their models would be "cost prohibitive." They say it would also require a dizzying degree of computational power — more than was required to train Claude in the first place.

"The economics will likely never make sense," cognitive scientist Gary Marcus argued in a recent **essay**. "Sky high valuation of companies like OpenAI and Microsoft are largely based on the notion that LLMs will, with continued scaling,

become artificial general intelligence,” he points out. But this is fantasy, Marcus says. “There is no principled solution to hallucinations in systems that traffic only in the statistics of language without explicit representation of facts and explicit tools to reason over those facts.”

Moreover, it’s still far from certain whether advanced AI can ever be profitable. The signs of an investment bubble were already emerging by late 2023. A year later, and industry financials continue to confound even the world’s most astute money managers.

A **report** from Goldman Sachs from this past June says companies worldwide are poised to spend US\$1 trillion on AI infrastructure in the near term, with little to show for it. And another **research note** issued by the bank in September adds some clarity. It says investment risk stems mostly from how “a handful of tech stocks account for an uncommonly high share of market capitalization.” The solution, according to one of the bank’s top strategists, is for investors to pivot toward “smaller technology companies and other parts of the market, including in the old economy, which will enjoy the growth of more infrastructure spend.”

Translation: the diverse global AI ecosystem will continue to evolve in tangible and productive ways. The technology’s benefits will be diffuse. But big tech behemoths are burning cash on the basis of thin promises.

“After years of pushing out increasingly sophisticated AI products at breakneck speed,” Bloomberg **reported** in mid-November, “three of the leading AI companies are now seeing diminishing returns from their costly efforts to build newer models.”

OpenAI and Google both released long-awaited new products in early December. OpenAI claims its **o1 model** — dubbed “Strawberry” and available to paid subscribers only — demonstrates reasoning capacity. The company also unveiled its AI video generator, Sora, to **mixed reviews**, and plans in early 2025 to release a new **o3 model**, which it believes is creeping toward AGI.

Google then revealed an upgraded Gemini model, which it envisions as a do-all personal aide. “I’ve dreamed about a universal digital assistant for a long, long time as a stepping stone on the part to artificial general intelligence,” Google

DeepMind CEO Hassabis **told** *WIRED* magazine. Anthropic has yet to release a new model due to disappointing progress.

Silicon Valley is indeed forging ahead. Industry researchers are focused on whether models' efficiency and accuracy can be boosted by allowing systems more time to process user inputs. Yet exploring this approach will require sustaining eye-watering levels of spending. "Tech companies like to make two grand pronouncements about the future of artificial intelligence," **says** journalist Matteo Wong. "First, the technology is going to usher in a revolution akin to the advent of fire, nuclear weapons, and the internet. And second, it is going to cost almost unfathomable sums of money."

One of those tech companies, Meta, spent upwards of **US\$40 billion** on new AI-dedicated data centres and hardware just in 2024. That's slightly more than the GDP of Estonia. Amazon spent around US\$75 billion — the economic output of Ghana. The cost of training the newest AI models alone now eclipses US\$1 billion. Anthropic's CEO Amodei thinks this might reach US\$100 billion by 2027. Elsewhere, Sam Altman estimates a US\$7 trillion expansion of the semiconductor industry is needed to achieve AGI. The OpenAI CEO is seeking to finance this quest by tapping the sovereign wealth funds of Gulf monarchies. This could seed a **geopolitical dilemma** down the line for the United States, in particular, and for liberal democracies in general.

There's also severe dissonance between how Silicon Valley ideologues promote the benevolent possibilities of AI and the ruthless tactics being used to develop it. Even if big tech companies can build AGI, there's little reassurance the technology will usher in a more equitable world.

"Over the past years, safety culture and processes have taken a back seat to shiny products," lamented Jan Leike in May 2024, when he was OpenAI's lead on aligning frontier AI systems with human values. The sentiment formed part of a lengthy **thread on X** announcing his resignation. A month later, more than a dozen current and former researchers from OpenAI, Google DeepMind and Anthropic claimed in an **open letter** that their bosses frequently vetoed their safety concerns. These same tech giants got to where they are through **harvesting training data without creators' consent**. The precarious class of **low-wage workers** fuelling the industry's boom are badly exploited. Antitrust laws are routinely defied; taxes and

finances go unpaid. Social media companies are also **dismantling content safeguards** despite vowing to rein in deepfakes and virulent false material. OpenAI exempts its top-tier users from visual watermarks on video content made using Sora.

“The real promise of AI is unlikely to become reality by itself,” Daron Acemoglu, a 2024 Nobel Prize-winning economist from the Massachusetts Institute of Technology, **wrote** recently. “It requires AI models to become more expert, better powered by higher-quality data, more reliable and more aligned with the existing knowledge and the information-processing capabilities of workers. None of this appears to be at the top of Big Tech’s agenda.”

What’s more, Silicon Valley’s ambassadors have embraced the cynical strategy of asking for regulations while pushing back against any proposed laws that have teeth.

OpenAI’s head of global policy suggested in a recent **interview** that the responsible use of AI will be determined by “whether democratic AI is going to prevail over autocratic AI.” And all major Western tech companies in July 2023 signed on to the Biden administration’s voluntary AI safeguards. A month earlier, Sam Altman himself called for the creation of a global AI-monitoring organization akin to the International Atomic Energy Agency. “We face existential risk,” Altman **told** an audience in the United Arab Emirates.

But these talking points mask big tech’s growing track record of ducking democratic checks and balances. In part, this is achieved through **colossal lobbying efforts**. For example, Gavin Newsom, the governor of California — home to Silicon Valley — in September shot down a legally binding AI safety bill overwhelmingly passed by state legislators, after a tech industry coalition launched a public relations blitz on it. The new law would have both compelled tech companies to conduct safety tests on the industry’s largest models and held companies liable for the harms they cause. The bill also called for tech companies to install a so-called kill switch to prevent systems from going rogue. Newsom authorized a smaller **patchwork set of laws** instead.

The second Trump administration’s approach to tech regulation hasn’t yet taken shape. But the president-elect’s instincts trend toward law-of-the-jungle-style corporate freedom. Media **reports** indicate Trump’s plan is to ditch President Biden’s executive order on AI safety, signed in October 2023, in favour of swift



deregulation. And the campaign period saw a raft of tech barons rush to back Trump. In particular, the list includes Elon Musk, who has been appointed to lead the president-elect's new Department of Government Efficiency, which is tasked with advising the White House on how to slash US federal expenditures. This despite Musk's companies collectively facing dozens of federal probes and lawsuits into alleged malpractices.

Taken together, Silicon Valley's timelines and ambitions for AGI point less to a solid business case than a techno-libertarian pursuit of feral entrepreneurialism. For years, this quest was epitomized by Mark Zuckerberg's reckless ethos of wanting to "move fast and break things." Lately, it has become just as apparent in Sam Altman's **goal** to automate away the "median human." Elon Musk, for his part, has contributed by helping **spread the idea** of government run by "high status males." In a self-published **manifesto** posted on his company's website, venture capitalist Marc Andreessen quotes Filippo Tommaso Marinetti, a twentieth-century Italian futurist and ally of fascist dictator Benito Mussolini: "There is no masterpiece that has not an aggressive character. Technology must be a violent assault on the forces of the unknown, to force them to bow before man." In a section titled "The Enemy," Andreessen lists the precautionary principle — the idea that high-risk endeavours should be executed slowly to avoid irreversible damage — as one to be denounced. Sustainability and social responsibility, too, he says, are "bad ideas" corrosive to human progress. Since Trump's re-election, Andreessen has become a key adviser to him on tech and economic policy.

Such hubris is also evident in Silicon Valley's insular culture and fascinations with **transhumanism** and **tech-enabled immortality**. Iconoclast investor Peter Thiel has been an avid proponent of the **seasteading movement**, which seeks to enable plutocrats to evade state authority by creating artificial floating communities in international waters. Meanwhile, the planet burns and nearly a third of humanity has still never accessed the internet.

"Can private companies pushing forward the frontier of a revolutionary new technology be expected to operate in the interests of both their shareholders and the wider world?" That question was posed earlier this year in a guest **essay** for *The Economist* by Helen Toner and Tasha McCauley, two former board members of OpenAI. Both women were jettisoned from the company in the power struggle that

erupted following Sam Altman being briefly fired in November 2023, allegedly for lying to the board about key decisions and internal safety protocols.

Altman was reinstated just five days after he was ousted, thanks to concerted pressure from OpenAI's industry allies and main investors. "Our particular story offers the broader lesson that society must not let the roll-out of AI be controlled solely by private tech companies," Toner and McCauley write. "Only through a healthy balance of market forces and prudent regulation can we reliably ensure that AI's evolution truly benefits all of humanity."

Even then, data constraints and sheer economics mean AI itself may never reach the heights its greatest enthusiasts fantasize about. It's more likely to "intensify and solidify the structure of the present," suggests Toronto-based writer Navneet Alang. "There will still be cracks in the sidewalk. The city in which I live will still be under construction. Traffic will probably still be a mess, even if the cars drive themselves."

AI is bound to alter the way humans live, work and interact. In ways good, bad and still unknowable — and almost certainly different from those envisioned by the Silicon Valley evangelists.