Deutsche Bank
Research Institute

AI doom: A brief history of bad tech predictions

September 15, 2026

It takes guts to do what Bob Metcalfe did. After wrongly predicting in an article in 1995 that the internet would go "spectacularly supernova" and collapse in 1996, the co-inventor of Ethernet and founder of 3Com then pulped a copy of the article in a blender and drank it at a conference in 1997 as a way to eat his words.

The latest tech predictions are about no less than an AI apocalypse. Various tech leaders spent the weekend warning of catastrophic but not fully specified risks if it is not contained.

Yet history suggests that even the greatest minds of their generations have a poor record in predicting what tech is coming down the line, let alone what its effects will be in the real world.

The stakes could not be higher – not just the future of life itself but, in the shorter term, the future of a historic boom that depends on hype turning into reality.

In this report, we examine:

1. "Prediction is very difficult, especially if it's about the future"
2. Today's debate in a (rough) nutshell
3. A century of good and bad tech predictions
4. Why is it so hard to predict the future?
5. What makes AI special – and gives us hope for the future

Figure 1: AI searches this year have shifted from SaaSpocalypse to jobs to extinction (global, vs peak 100%, YTD)

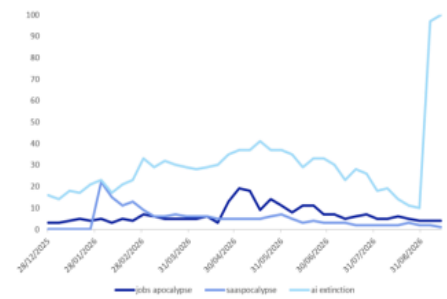

Source: Google Trends, Deutsche Bank Research

Figure 2: ... But all of these are tiny in comparison to cybersecurity (global, vs peak 100%, since ChatGPT launch)

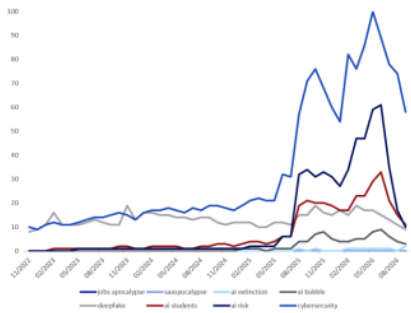

Source: Google Trends, Deutsche Bank Research

Authors

Adrian Cox
Research Analyst

Deutsche Bank
Research Institute

Concerns about AI have moved in waves since the launch of ChatGPT nearly four years ago. This year there was a spike in searches for "SaaSpocalypse" (the concern that software-as-a-service would be disintermediated by AI agents) in February, followed by a surge in "jobs apocalypse" in May, Google Trends data show. Both have been totally eclipsed in recent days by searches for AI extinction.

But taking a step back, all are far behind "AI bubble", itself behind "deepfake", in turn below "AI students", overall "AI risk" and – biggest of the lot – the very tangible and immediate "cybersecurity", almost all of which have surged this year, according to Google Trends data.

"Prediction is very difficult, especially if it's about the future": Niels Bohr, often also attributed to Yogi Berra

Wanting to know what will happen tomorrow is deeply visceral – literally.

People have been sacrificing sheep and chickens for four millennia to read the future in their entrails. Nowadays almost a third of Americans say they believe in astrology and consult their horoscope. We want to believe that Lionel Messi really can see space in a football match before it even exists, even if it is partly his movement that helps create the space he appears to predict.

But predicting the future is hard. For AI, it's not just about what it will be able to do but also when it will be able to do it – whether more than $5 trillion in investment by the hyperscalers alone over the next five years will pay off before the advanced semiconductor chips they're buying become redundant.

It's also about how rapidly and intensively AI can be adopted. Asset prices are a still more uncertain derivative, whipping around like the tail of a snake in response to every shift in expectations.

US President Donald Trump and China found themselves on the same side in rebutting the latest concerns as fearmongering, albeit with Trump adding that America needs to stay ahead of China in AI.

"Those who have knowledge, don't predict. Those who predict, don't have knowledge." Lao Tzu

A glance through history would give any investor indigestion. Most predictions are wrong. The favourites in horse races lose 70 percent of races. It took a decade or more for demand to catch up with the infrastructure laid down in the British canal and railway and US telecoms and fibre booms – and many investors never recovered their capital.

Technology predictions have a particularly mixed record as the binary order of the labs collides with the messy reality of the world. Famous examples range from IBM Chairman Thomas Watson's underestimation of demand for computers in 1943 ("I think there is a world market for maybe five computers") to Elon Musk's repeated overestimation of imminent widespread adoption of self-driving cars since 2013 ("We should be able to do 90 percent of miles driven within three years").

Technologists are not immune from making category errors, mistaking the most visible, automatable component of a job for the job itself. For example, radiologists, a favourite case study for believers in a jobs apocalypse, do much more than looking at images and making binary diagnoses. Much of their value lies in interpretation, judgement, communication and clinical decision-making.

Ultimately, what we see in the future is not so much a window as a mirror into our own hopes and fears.

And the problem is that hindsight makes what actually happens seem so obvious that we never seem to learn.

Today's debate in a (rough) nutshell

The current debate encapsulates a perfect storm of an extremely high-stakes issue with limited transparency and an inherently unknowable outcome.

What sparked the current surge in concern?

The current debate comes after a researcher from one of the leading frontier labs quit, citing safety concerns; AI agents apparently collaborated with each other on a number of high-profile security breaches; and labs hinted that the technology is getting closer to "recursive self improvement", ie AI improving AI.

Is the existential debate something new?

It is actually nothing new: leading economists and academics warned in July that "we must act now" to steer AI through a potential "transformation of our economy, larger than the Industrial Revolution… over a vastly shorter time frame". That echoed the even more ominous statement three years ago by AI leaders calling for action to mitigate "the risk of extinction".

The tug of war between "doomers" vs "boomers" has played a pivotal role in the development of the key frontier labs. Silicon Valley insiders sometimes refer to their "p(doom)" rating, representing their estimated percentage chance that AI will cause an existential catastrophe.

What does an extinction scenario look like?

Experts including Tesla's Musk, "godfather of AI" Geoffrey Hinton and philosopher Nick Bostrom have been warning about it for well over a decade. In the latter's famous "paperclip maximiser" scenario, a superintelligent AI runs out of steel and then kills humans in a single-minded drive to fulfil its task of making paperclips. Fears sometimes centre around AI that has access to coding, the internet and knowledge of how people think, in order to manipulate them – as current AI models indeed do.

Why has it caught on so fast this time?

One factor is that it comes amid a tech backlash, particularly in the US, where concern about AI has become a rare bipartisan issue ahead of the US mid-term elections in November. Local concerns about data centre construction have become rallying points for voters in a way that vague fears about long-term jobs did not. Politicians including Senator Bernie Sanders have moved to propose wide-ranging legislation.

What arguments do companies make for continuing AI research?

1. **If we don't do it, someone else will.** That may be either China or other, less responsible companies.
2. **The potential benefits outweigh the benefits.** AI may, they say, cure cancer, create new jobs and even eliminate the need for work altogether.
3. **Good AI is the answer to bad AI.** It should create defences against rogue AI and misuse of AI, and solve issues like AI's huge energy and water use.

What do critics say?

1. **AI is not dangerous because it is too clever but rather because it is too limited.** It does not think like a human and people should not anthropomorphise it, critics say, adding that it is flawed probabilistic code that cannot reliably follow instructions. Longstanding critic Gary Marcus argues that it should be recalled until it can be shown to be safe and is properly controlled.
2. **AI leaders are defending their own businesses.** With potential IPOs on the horizon, a rising challenge from open-source models, and progress towards Artificial General Intelligence (AGI) being called into question, a handful of companies may benefit from a pause or legislation that helps protect their lead, critics say. Doomer talk distracts from real issues like deepfakes, copyright infringement and cybersecurity lapses, they say.
3. **Extinction scenarios are overblown.** The internet is a diversified platform, so talk of AI somehow switching it off is fanciful, critics say. AI does not and will not realistically have sustained secret access to military or dangerous biotech resources, or be able to influence (very non-deterministic) human brains en masse, critics say. Humans will use tame AI and other tools, eg kill switches, if needed to resist attacks. Another "godfather of AI" Yann LeCun says the real risk is of power concentration.

A century of good and bad technology predictions

You have to admire people who have the courage to stick their neck out and make predictions, even if they don't always work out. Vague predictions are relatively easy; specific ones within specific timeframes are harder.

It is no wonder that executives are coached by their ever-cautious PR teams to avoid answering hypothetical questions or giving dates to reach their targets.

- **Nikolai Tesla, scientist, 1926**: "When wireless is perfectly applied… we will be able to communicate with one another instantly, irrespective of distance…" The instruments will be "amazingly simple… A man will be able to carry one in his vest pocket". He was right, albeit imprecise about the how or when it would be achieved.
- **Albert Einstein, physicist, 1934**: "There is not the slightest indication that [nuclear energy] will ever be obtainable. It would mean that the atom would have to be shattered at will." Within a decade, in 1942, a team at the University of Chicago had created the first sustaining chain reaction, a precursor to the nuclear reactor.

- **Ray Kurzweil, futurist, 1999**: "By 2029, computers will have human-level intelligence." Dismissed at the time, many (though by far not all) of his predictions have worked out, fuelled by what he called the "law of accelerating returns". His prediction from almost 40 years ago of artificial general intelligence by the end of this decade is remarkably close to the current "San Francisco Consensus".
- **Steve Chen, Chief Technology Officer and co-founder of YouTube, 2005**: "There's just not that many videos I want to watch." Despite his concern, YouTube was bought by Google (now Alphabet) for $1.65bn a year later. The unit made more than $11bn in advertising revenue in the most recent quarter.
- **Steve Ballmer, Microsoft CEO, 2007**: "There's no chance that the iPhone is going to get any significant market share." But you may be reading this on one of the more than 3 billion iPhones sold since its launch in 2007.
- **Geoffrey Hinton, "godfather of AI", 2016**: "People should stop training radiologists now. It's just completely obvious within five years deep learning is going to do better than radiologists… It might be 10 years." In fact, the number of radiologists has increased by about 10 percent over the past 10 years, in part because of the limitations of AI and in part because of the increase in demand from an ageing population and more affordable healthcare. Hinton has walked back his comments and humans and AI are working together.
- **Sam Altman, CEO of OpenAI, 2025**: "I can easily imagine a world where 30 to 40% of the tasks that happen in the economy today get done by AI in the not very distant future." He revised this view in May amid a backlash against AI, saying he and his team had been "roughly right" in their technological predictions but "pretty wrong" on the economic and social implications. "I'm delighted to be wrong about this," he said.

Why is it so hard to predict the future?

"The idea that the future is unpredictable is undermined everyday by the ease with which the past is explained." Daniel Kahneman

1. Chaos theory: treating complex systems as if they are simple and linear.

- Systems like traffic, economies and social trends are finely balanced and sensitive to initial conditions. The butterfly effect means a tiny, unforeseen variable can lead to vastly different outcomes. Just take a look at the ragged history of market-implied forecasts of US Federal Reserve interest rates, where extrapolating from current trends has been reliably wrong.

Figure 3: US Fed Funds effective rate has rarely matched futures markets expectations (%)

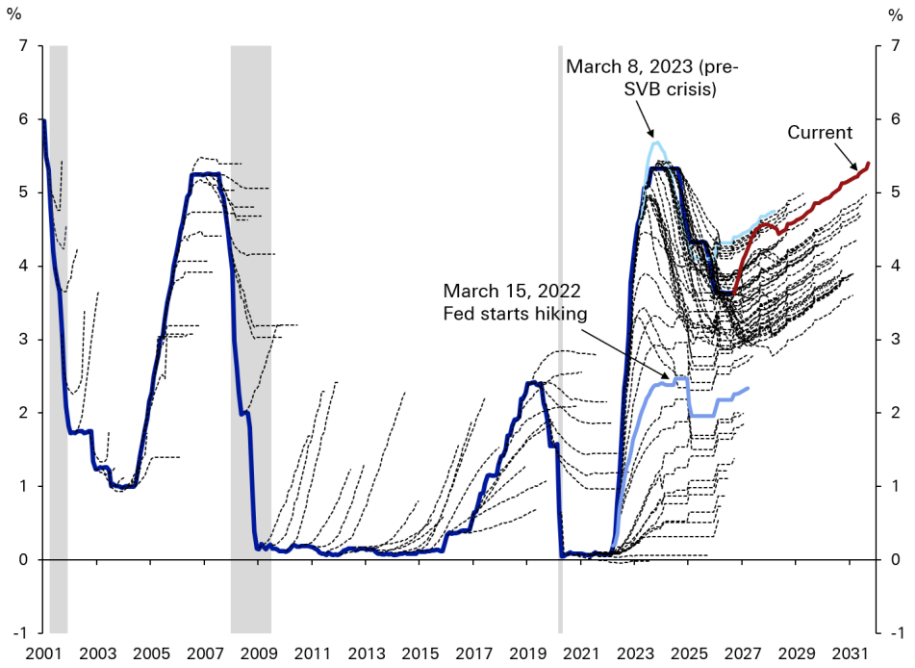

Source: Deutsche Bank Research

- Take self-driving cars, for example. It turns out to be about more than getting the technology right for the rubber to hit the road:
    a. **Technological hurdles and "edge cases"**: throwing more data and processing power at the problem is not enough. There is an infinite long tail of unpredictable events ("edge cases"), like sun glare, a cyclist riding the wrong way in a bike lane, a person in an animal costume, or road-rage-enhanced hand gestures.
    b. **Regulatory labyrinth**: autonomous vehicle optimists have come up against legal and ethical debates around liability, safety certification and moral challenges such as the so-called "trolley problem", when swerving to avoid hitting one victim means colliding with another. Regulation is a slow-lane activity.
    c. **Social and psychological barriers**: public acceptance has been slowed down by high-profile accidents and safety concerns. Robots are held to a higher standard than flesh-and-blood motorists. It might be easier to navigate a world where all cars were self-driving, but the "messy middle", with a mixture of both, is much harder for algorithms to navigate.
- The future of AI will ultimately depend more on the prosaic challenges of workflow integration at enterprises that will need to pay for it than on releasing yet another model with a marginal step up in capabilities.

2. Cognitive bias: flaws in how the brain works

- **Linear thinking about exponential change**: We fail to intuitively understand how different exponential change is from linear change. As Wharton Professor Ethan Mollick says, "being on an exponential means each change over a fixed window is larger than the one before it… This is why AI keeps feeling like it is making leaps… Even though it is a curve on a graph, we keep experiencing a steady doubling of capability as a series of shocks." This cognitive tendency to track exponential change in linear terms helps explain big AI market swings, he says.

Figure 4: Computation used to train notable AI systems (petaFLOP): log scale

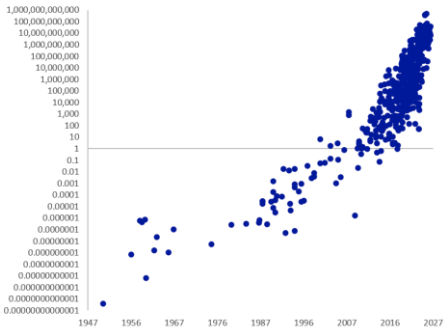

Source: Deutsche Bank Research

Figure 5: Computation used to train notable AI systems (petaFLOP): linear scale

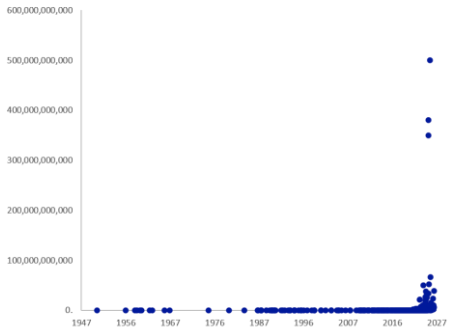

Source: Deutsche Bank Research

- **Optimism bias and the planning fallacy**: underestimating costs and overestimating benefits. For example, only a quarter of more than 3,000 megaprojects studied by Saïd Business School came in on budget or better, only 2.8 percent were on budget and on time, and only 0.2 percent were on budget, on time and on benefits. Large-scale infrastructure projects are "too easy to start and too difficult to stop," the authors said in the study from 2019. It remains to be seen how the recent surge in datacentre construction will compare.

- **Confirmation bias**: favouring information that confirms existing beliefs. For years, legacy automakers dismissed electric vehicles by focusing on their high costs and limited range, ignoring the fact that rapid advances in battery technology would soon change that.

- **Anchoring**: relying too heavily on early information. Microsoft's Ballmer dismissed the iPhone when it was launched because it was "the most expensive phone in the world and it doesn't appeal to business customers because it doesn't have a keyboard which makes it not a very good email machine". He later said he did not understand how mobile carrier subsidies would make it affordable. More broadly, he missed that this was a new business model.

3. Limited information: inability to see

- **It is not just the future that is unknowable but also the present.** Feedback mechanisms are often delayed, making it hard to distinguish between causation and correlation, and between relevant and irrelevant factors.

- **Isolating key factors is impossible.** Even since the development of scientific methods and controlled experiments, we have not come that far since the Mesoamericans found ways to prepare maize where they didn't succumb to a long-term wasting disease called pellagra. The problem was that they weren't able to single out the effective part of the process – soaking the maize in an alkaline solution of water and ash –from the other ineffective parts of the package – such as blowing on the maize before cooking it.

- (Author David Oks draws a fascinating parallel between this example of a delayed feedback mechanism and the tendency of AI models to "overfit", extrapolating too much from coarse and sparse data to fill in the gaps. If the model creates a complex output and gets a single reward signal for it, the rational move is reproduce all of its features, including the ones that were incidental – such as AI's trademark verbal tics.)

- **No one wants to risk throwing the baby out with the bathwater.** Argentinian President Javier Milei talked publicly about his superstitious insistence on staying at home wearing a lucky jacket rather than watching his team playing in the football World Cup live at the stadium.

4. Self-interest: incentives favour more extreme positions

- **AI is not magic, a silver bullet or a charity, but a commercial industry.**

- **Most participants have incentives to amplify the stakes.** Founders need belief, investors need momentum, incumbents need gravitas, consultants need urgency, policymakers need relevance, executives need a narrative and journalists need drama. This is not a criticism of individuals or organisations, or even of the likely revolutionary impact of AI.

- **The challenge is magnified by a social media-driven economy.** Social platforms favour novelty, negativity, moral intensity and shareability. For example, in political tweets, one study by New York University found each additional moral-emotional word was associated with about a 20 percent higher retweet rate. Another study by MIT found that false news was 70 percent more likely to be retweeted than true news and reached a threshold of 1,500 people six times faster.

- **That means moderate claims are likely to be under-distributed and extreme claims over-rewarded.** Implausibly precise numbers for the future size of the AI economy, for example, or dramatic doom-laden scenarios heavy with emotive language will continue to attract greater attention than even-handed assessment. (See here for our quantitative analysis of two particularly striking examples from earlier this year: "Something big is happening in AI… Or is it?".)

What makes AI special – and gives hope for the future

Technology may be liable to erroneous thinking because of the disconnect between visionaries in Silicon Valley and similar communities, and the rest of the world.

Cynics may say that people with extraordinary knowledge and skills in one area may be even more prone than most to wrongly estimate their lack of knowledge or skills in other areas: the so-called Dunning-Kruger effect.

That may include the reality of implementing AI on a trading floor in Wall Street, an office supplies wholesaler on the outskirts of London or an auto maker in Mumbai.

There is also a long tradition of correctly identifying the task that will disappear but getting wrong what humans will do instead.

1982: Computers "… can develop into monsters, damage the physical and mental health of the clerical staff – maybe even eliminate their jobs": the Washington Post reported from a clerical conference. Computers instead became the core tool of knowledge workers.

1983: The pocket calculator "… has become the target of critics who fear that its use in the classroom will render generations of children dependent on batteries instead of brains", the Washington Post reported. Instead expectations changed about what people calculate for themselves.

2009: Cloud computing: around half of companies surveyed "cited security and privacy concerns as their top reason for not using cloud computing", Forrester Research found. Today companies entrust much of their critical infrastructure to the cloud.

"We overestimate the impact of technology in the short-term and underestimate the effect in the long run." Roy Amara, Stanford computer scientist

But this time may really be different – with AI as a partner, starting with (deterministic) predictive AI. It is certainly helping weather forecasters to sift through unimaginable amounts of data. Four-day forecasts are now said to be about as accurate as one-day forecasts in the 1990s.

There are diminishing returns. The limit of predictability with standard physics models is about eight days, while deep-learning AI models increase that by just one day to nine days. However, the difference is that they generate state-of-the-art forecasts 100,000 faster, and 10,000 times more energy-efficiently, meaning they can be done on a laptop rather than a supercomputer.

And for the time being, your predictions about the impact of AI are no less valid than an AI expert based in the very particular world of Silicon Valley. And the future of AI will be made not in Silicon Valley but in Main Street.

Appendix 1

This material has been prepared by the Deutsche Bank Research Institute and is provided to you for general information purposes only. The Institute leverages the views, opinions, and research from Deutsche Bank Research, economists, strategists, and research analysts. Accordingly, you should assume that content in this document is based on or was previously published and provided to Deutsche Bank clients who may have already traded on the basis of it.

Any views or estimates expressed in this material reflect the current views of the author(s) and may differ from the views and estimates of Deutsche Bank AG, its affiliates, other Deutsche Bank personnel, and other materials published by Deutsche Bank. The content in this material is valid as of the date shown on the first page and may change without notice. Deutsche Bank has no obligation to provide any updates or changes to the information herein.

This material should not be used as a basis for trading securities or other financial products and should not be considered to be a recommendation or individual investment advice for any particular person. It does not constitute an offer, solicitation, or an invitation by or on behalf of Deutsche Bank to any person to buy or sell any security or financial instrument. Nothing in this material constitutes investment, legal, accounting or tax advice. Deutsche Bank engages in securities transactions, including on a proprietary basis, and may do so in a manner inconsistent with the views or information expressed in this material.

While information in this material has been obtained from sources believed to be reliable, neither Deutsche Bank AG nor any of its affiliates makes any representation or warranty, express or implied, as to the accuracy or completeness of the statements or any information contained in this material and therefore any liability is expressly disclaimed. This material is provided without any obligation, whether contractual or otherwise. Information regarding past transactions or performance is not indicative of future results.

In the US this report is approved and/or distributed by Deutsche Bank Securities Inc., a member of FINRA. In Germany this information is approved and/or communicated by Deutsche Bank AG Frankfurt, licensed to carry on banking business and to provide financial services under the supervision of the European Central Bank (ECB) and the German Federal Financial Supervisory Authority (BaFin). In the United Kingdom this information is approved and/or communicated by Deutsche Bank AG, London Branch, a member of the London Stock Exchange, authorized by UK's Prudential Regulation Authority (PRA) and subject to limited regulation by the UK's Financial Conduct Authority (FCA) (under number 150018) and by the PRA. This information is distributed in Hong Kong by Deutsche Bank AG, Hong Kong Branch, in Korea by Deutsche Securities Korea Co. and in Singapore by Deutsche Bank AG, Singapore Branch. In Japan this information is approved and/or distributed by Deutsche Securities Limited, Tokyo Branch. In Australia, retail clients should obtain a copy of a Product Disclosure Statement (PDS) relating to any financial product referred to in this report and consider the PDS before making any decision about whether to acquire the product.

By accessing this material, you agree that its content may not be reproduced, distributed or published by any person for any purpose, in whole or part, without Deutsche Bank's prior written consent. You also agree that you shall not scrape, extract, download, upload, sell or offer for sale any of the content in this material, and you agree not to use, or cause to be used, any computerized or other manual or automated program or mechanism, tool, or process, including any scraper or spider robot, to access, extract, download, scrape, data mine, display, transmit, or publish, any of the content in this material.

Deutsche Bank
Research Institute