

Advanced AI

Beyond the AI hype - the FAQ

Common Questions Around AI's Transformative Potential

This analysis addresses the most common questions about AI's transformative potential, drawing on insights from policymakers, experts, and research across the ecosystem. It offers a concise, non-technical summary to help policymakers and their teams cut through the noise and engage with the key issues. Designed as a brief and accessible entry point into the AI debate, it complements our full report, [Beyond the AI Hype: A Critical Assessment of AI's Transformative Potential](#).

Table of Contents

Introduction.....	2
Are current valuations and investments completely unsustainable, or the start of a trillion-dollar revolution?.....	3
Is AI actually improving or just getting better at taking tests?.....	7
AI still makes basic mistakes, surely that will limit its potential for transformation?.....	11
We've seen tech hype cycles before—why is this time different?.....	14
Is AI hitting fundamental limits or will new methods unlock the next leap?.....	17
Is transformative AI impact decades away or is change coming much faster than expected?	20
Acknowledgements.....	26
Contact us.....	26
Reference List.....	27

Introduction

Artificial intelligence (AI) dominates headlines, boardroom discussions, and policy debates, with claims about its potential ranging from revolutionary breakthrough to overhyped disappointment. This polarised landscape makes it challenging to separate genuine progress from personal belief or marketing buzz. Sceptical voices offer important counterpoints to the most enthusiastic AI predictions, highlighting limitations, failures, and historical parallels to previous technological bubbles. Some of these concerns are well-founded and point to real challenges that need addressing. Others, however, may underestimate the unique characteristics of current AI progress and the evidence in favour of transformative impact.

This analysis examines the most common questions about the transformative potential of AI, based on our interactions with policymakers, colleagues, academic papers, industry analyses, domain experts and others in the policymaking ecosystem. We seek to evaluate each against available evidence and present a non-technical summary for general audiences.

Unpacking the signal from the noise is critical for ongoing discussions around the governance and development of AI, and it is particularly important that this topic be easy to understand for non-technical audiences. The aim of this document is to provide a primer for policymakers and their staff, a baseline assessment of that will serve as a solid entry point into the AI debate.

Our assessment suggests that while some scepticism about AI is warranted, many dismissive arguments underestimate AI's unique characteristics and potential for transformative impact within years rather than decades. This recognition should inform responsible action, not reckless acceleration. The trajectory of AI development is not guaranteed to be positive by default and presents significant challenges such as uneven distribution of benefits, potential power concentration, bias concerns, mounting energy demands, and emerging systemic risks as AI integrates with critical infrastructure. The ultimate impact of AI remains in our hands, guided by the decisions we make about how this powerful technology should serve humanity's diverse needs and values.

This document is intentionally brief and focused—for a deeper analysis, please refer to our comprehensive report "[Beyond the AI Hype: A Critical Assessment of AI's Transformative Potential](#)".¹ For this analysis to provide a timely snapshot, we drew upon all available sources of information, preferentially relying on the work of established organisations and focusing on long-term trends and significant developments. However, given the rapid advancement of the AI field, there is often limited academic literature or in-depth analyses available, requiring us to sometimes draw conclusions from fewer datapoints than ideal. Consequently, some specific details may become dated. We encourage readers to reach out to us directly for our most current perspectives or to inquire about an arrangement of updated briefings for their teams and policy units.

¹Janků, D., Reddel, F., Graabak, J. and Reddel, M., 'Beyond the AI Hype – A Critical Assessment of AI's Transformative Potential', Centre for Future Generations, 2025, available at: <https://cfg.eu/beyond-the-ai-hype>.

Are current valuations and investments completely unsustainable, or the start of a trillion-dollar revolution?

Widespread doubts about AI often center on financial issues, with observers pointing to sky-high valuations and minimal revenue reminiscent of past tech bubbles. However, crucial differences from previous tech booms exist: major AI providers already generate substantial revenue, corporate spending is backed by tangible demand, and rapid cost decreases make the technology increasingly accessible. While some individual startups might be overvalued, the wider AI ecosystem shows strong fundamentals through demonstrated returns, proven infrastructure investments, and rapid technological progress, suggesting this boom is built on more than mere speculation.

Some of the loudest scepticism surrounding AI right now is focused on the financials. Here are some [key aspects](#)² observers warn about based on previous tech bubbles:

- **Soaring valuations with minimal revenue.** The revenue-to-enterprise value ratio for all Software (System & Application) is [about 1:12](#)³ and for [financial services 1:19](#)⁴, while AI startups often fetch billion-dollar valuations with less favourable [ratios of around 1:30](#)⁵, although some reach 1:150. Even OpenAI, with real substantial income, stands [at about 1:46](#)⁶. Higher ratios compared to industry standard [might suggest](#) that the company is overvalued, or that investors believe that future sales will greatly increase⁷.
- **Lack of monetisation strategies.** Adding to the concerns, some AI firms still lack clear monetisation strategies; many are acquired primarily for their talent (“acqui hires”)—for instance, [CharacterAI by Google](#)⁸, [AdeptAI by Amazon](#)⁹, or [InflectionAI by Microsoft](#)¹⁰.

²Floridi, L., 'Why the AI Hype is Another Tech Bubble', Philosophy & Technology, Springer, 2024, available at <https://link.springer.com/article/10.1007/s13347-024-00817-w>.

³NYU Stern, 'Price to Sales Ratios', 2025,

https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/psdata.html.

⁴NYU Stern, 'Price to Sales Ratios', 2025,

https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/psdata.html.

⁵Ansari, S., 'AI Valuation Multiples 2024', Aventis Advisors, 2025, available at

<https://aventis-advisors.com/ai-valuation-multiples/>

⁶Ansari, S., 'AI Valuation Multiples 2024', Aventis Advisors, 2025, available at

<https://aventis-advisors.com/ai-valuation-multiples/>

⁷Hayes, A., 'What Is Enterprise Value-to-Sales (EV/Sales)? How to Calculate', Investopedia, 2021, available at <https://www.investopedia.com/terms/e/enterprisevaluesales.asp>

⁸The Wall Street Journal, 'Struggling AI Startups Look for a Bailout From Big Tech', 2024, available at <https://www.wsj.com/tech/ai/struggling-ai-startups-look-for-a-bailout-from-big-tech-3e635927>

⁹The Wall Street Journal, 'Struggling AI Startups Look for a Bailout From Big Tech', 2024, available at <https://www.wsj.com/tech/ai/struggling-ai-startups-look-for-a-bailout-from-big-tech-3e635927>

¹⁰Fast Company, 'Microsoft's Inflection AI grab cost more than \$1 billion', 2024, available at <https://www.fastcompany.com/91069182/microsoft-inflection-ai-exclusive>

- **AI Washing.** Many businesses now push an AI angle or adopt “AI” signifiers despite having little or no actual connection with the technology. This [“AI washing”¹¹ pays off with investors¹²](#), as past studies show, but are a sign of a fad taking over the market.

Yet there are crucial differences from the dot-com days. Major AI providers already generate sizable revenue: Microsoft [surpassed \\$10 billion¹³](#) in AI revenue in 2024, while OpenAI [expects \\$3.7 billion¹⁴](#), and Amazon and Alphabet (the parent company of Google) also [report multibillion-dollar¹⁵](#) AI income. These aren’t speculative forecasts—they’re real earnings. A McKinsey Global Survey from May 2024 shows that organisations deploying generative AI report tangible value: on average, [39% of surveyed people report cost reductions and 44% report increased revenue¹⁶](#).

And this is being backed by investments into the infrastructure needed to keep scaling AI. Major tech firms are pouring in [hundreds of billions¹⁷](#) into this technology. Partly, this is because they see synergies with existing income streams, most notably cloud computing, but they are also responding to tangible demand: providers report [more interest in AI solutions than their current capacities can handle¹⁸](#).

On top of that, AI’s costs are dropping at breathtaking speed. OpenAI’s Dane Vahey notes that AI model [prices fell by 99% in just 18 months¹⁹](#)—an amount of decline for which [solar power needed almost](#)

¹¹Wikipedia, 'AI washing', 2025, available at https://en.wikipedia.org/wiki/AI_washing

¹²Olson, P., 'Nearly Half Of All 'AI Startups' Are Cashing In On Hype', Forbes, 4 March 2019, available at <https://www.forbes.com/sites/parmyolson/2019/03/04/nearly-half-of-all-ai-startups-are-cashing-in-on-hype/>

¹³Seeking Alpha, 'Microsoft Corporation (MSFT) Q1 2025 Earnings Call Transcript', Seeking Alpha, 2025, available at <https://seekingalpha.com/article/4731223-microsoft-corporation-msft-q1-2025-earnings-call-transcript>.

¹⁴Hu, K. and Cai, K., 'OpenAI offers one investor a sweetener that no others are getting', Reuters, 28 September 2024, available at <https://www.reuters.com/technology/artificial-intelligence/openai-sees-116-billion-revenue-next-year-offers-thrive-chance-invest-again-2025-2024-09-28/>

¹⁵Forbes, 'AI Spending to Exceed a Quarter Trillion Next Year', Forbes, 2024, available at <https://www.forbes.com/sites/bethkindig/2024/11/14/ai-spending-to-exceed-a-quarter-trillion-next-year/?ref=platformer.news>.

¹⁶McKinsey, 'The State of AI in 2024', McKinsey & Company, 2024, available at <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

¹⁷Kindig, B., 'AI Spending to Exceed a Quarter Trillion Next Year', Forbes, 2024, available at <https://www.forbes.com/sites/bethkindig/2024/11/14/ai-spending-to-exceed-a-quarter-trillion-next-year/>.

¹⁸Kindig, B., 'AI Spending to Exceed a Quarter Trillion Next Year', Forbes, 2024, available at <https://www.forbes.com/sites/bethkindig/2024/11/14/ai-spending-to-exceed-a-quarter-trillion-next-year/>.

¹⁹Tsarathustra, 'Dane Vahey of OpenAI Says the Cost...', X, 2024, available at <https://x.com/tsarnick/status/1838401544557072751>.

[40 years](#)²⁰ and [many other technologies took even longer](#)²¹. If this rapid cost depreciation is even close to accurate (as shown in Figure 1 below), it suggests that while initial investments are very high, the technology is becoming dramatically cheaper on a per unit basis over time. The entry of DeepSeek early in 2025 showcased that further efficiency gains will likely make the AI sector more competitive in the future; as upfront capital costs reduce, so does the barrier to entry.

In conclusion, while some AI startups might be chasing hype with inflated valuations and shaky business models, the wider AI ecosystem rests on more solid foundations. Unlike the speculative dot-com era, today's leading AI companies are already delivering substantial revenue, major cost efficiencies, and tangible market demand. However, despite rapid progress and promising early returns, the future remains inherently unpredictable—AI is not a guaranteed money machine, and extraordinary gains many investors expect are not assured. While individual AI startups may be risky bets, the sector's demonstrated returns, solid infrastructure investment, and rapid technological progress suggest this boom is built on more than just speculation.




Figure 1: Comparison of the decline in costs for solar power (orange, in \$/Watt), DNA sequencing (blue, in \$/full human genome), and AI inference (green, in \$/1M tokens). The starting points for each technology are somewhat artificial but were selected to enable the 99% reduction comparison while

²⁰Our World in Data, 'Solar PV Prices', Our World in Data, 2024, available at <https://ourworldindata.org/grapher/solar-pv-prices>.

²¹'Data Page: The cost of 66 different technologies over time', Our World in Data, 2025, data adapted from Farmer and Lafond, available at <https://ourworldindata.org/grapher/costs-of-66-different-technologies-over-time>

capturing the steepest cost declines. Data sources: Solar power: [Our World in Data – Solar PV Prices](#)²², DNA sequencing: [Our World in Data – Cost of Sequencing a Human Genome](#)²³, AI inference: OpenAI's Dane Vahey's presentation recording shared by [Tsarathustra on X](#)²⁴, data can also be checked on [OpenAI's pricing page history](#)²⁵

Is AI actually improving or just getting better at taking tests?

A fundamental issue in assessing AI development is how to differentiate true capabilities from systems merely optimized to perform well on specific tests. While historical examples show benchmarks being "solved" without true understanding, newer evaluation methods and real-world performance metrics suggest more robust progress, with the caveat that we are still assessing an emerging field. AI is increasingly validated through direct comparisons with human experts in professional settings, showing genuine capabilities in medical diagnosis, legal work, and scientific research, with measurable business impact and growing adoption rates indicating value beyond artificial benchmarks.

This criticism targets a core challenge in evaluating AI: how to distinguish genuine capabilities from mere optimisation for a given benchmark. Some argue that celebrated AI benchmarks can be [poor measures of progress](#)²⁶ because companies can tailor models to the tests—with [direct or indirect access to test data](#)²⁷, fine-tuning for specific tests, and selection bias in reported results. As the authors of the book *AI Snake Oil* note, this [might have happened](#)²⁸ for example with GPT-4 boasting high scores on bar exams or coding challenges, only to find that it likely had those questions in its training data or figured out how to game the test format. Moreover, mastering artificial benchmarks may not translate to real-world success, where problems are messier and less clearly defined. As the above authors continue to argue, having an AI score in the 90th percentile on a lawyer's exam doesn't prove it can actually practice law; after all, "it's not like a lawyer's job is to answer bar exam questions all day."

This scepticism draws on historical AI examples. Benchmarks once seen as meaningful have been "solved" by models [lacking the underlying reasoning](#)²⁹ they were meant to test. For example, early

²²Our World in Data, 'Solar PV Prices', Our World in Data, 2024, available at <https://ourworldindata.org/grapher/solar-pv-prices>.

²³Our World in Data, 'Cost of Sequencing a Full Human Genome', Our World in Data, 2023, available at <https://ourworldindata.org/grapher/cost-of-sequencing-a-full-human-genome>.

²⁴Tsarathustra, 'Dane Vahey of OpenAI Says the Cost...', X, 2024, available at <https://x.com/tsarnick/status/1838401544557072751>.

²⁵OpenAI, 'Pricing', OpenAI, 2023, available at <https://web.archive.org/web/20230401144811/https://openai.com/pricing>.

²⁶Eriksson, M. et al., 'Can We Trust AI Benchmarks? An Interdisciplinary Review of Current Issues in AI Evaluation', arXiv, 2025, available at <https://arxiv.org/abs/2502.06559>.

²⁷Fodor, J., 'Line Goes Up? Inherent Limitations of Benchmarks for Evaluating Large Language Models', arXiv, 2025, available at <https://arxiv.org/abs/2502.14318>

²⁸AI Snake Oil, 'GPT-4 and Professional Benchmarks: The Wrong Answer to the Wrong Question', AI Snake Oil, 2023, available at <https://www.aisnakeoil.com/p/gpt-4-and-professional-benchmarks>.

²⁹Pacchiardi, L., Tesic, M., Cheke, L.G., and Hernández-Orallo, J., 'Leaving the barn door open for Clever Hans: Simple features predict LLM benchmark answers', arXiv, 2024, available at <https://arxiv.org/abs/2410.11672>

language models scored highly on [reading comprehension](#)³⁰ yet still failed at basic common-sense reasoning. There's also a valid concern about the artificial nature of many benchmarks—they often test capabilities in isolation, while real-world applications require handling multiple challenges simultaneously in less structured environments.

However, some new benchmarks suggest we're entering a more robust era. Take the ARC-AGI test, designed by AI sceptic François Chollet to [gauge reasoning akin to general intelligence](#)³¹. OpenAI's latest models [scored 75.7%](#)³²—surpassing [human performance of 64.2%](#)³³. In other words, while earlier AI models were good at recognising patterns, OpenAI's models can now solve problems that were specifically designed to be 'AI-resistant'—tasks that require the kind of fluid, adaptable reasoning that many experts thought would remain uniquely human for much longer. Even Chollet acknowledged genuine progress, commenting: "[When I see these results, I need to switch my worldview about what AI can do and what it is capable of.](#)"³⁴

More tellingly, direct comparisons between AI and human experts are moving beyond artificial tests. Take medical diagnosis: in a [controlled](#) study in diagnosing conditions from case reports, doctors working unaided reached 74% accuracy. With AI assistance, they improved slightly to 76%. But AI on its own? It [outperformed them both, achieving 90% accuracy](#)³⁵. Furthermore, in the largest randomised controlled trial of medical AI to date, cancer detection rates [rose by 29%](#)³⁶ with no increase in false positives while reducing radiologists' workload. Similarly, a Harvard Business School study found that consultants with access to GPT-4 [boosted productivity by 12.2%, worked 25.1% faster, and delivered 40% higher quality](#) than those without AI access³⁷. These aren't artificial benchmarks but real-world tasks performed by experienced professionals.

³⁰Pavlus, J., 'Machines Beat Humans on a Reading Test. But Do They Understand?', Quanta Magazine, 17 October 2019, available at <https://www.quantamagazine.org/machines-beat-humans-on-a-reading-test-but-do-they-underst-and-20191017/>

³¹ARC Prize, 'ARC Prize – What Is ARC-AGI?', ARC Prize, available at <https://arcprize.org/arc>.

³²'OpenAI o3 Breakthrough High Score On ARC-AGI-pub', ARC Prize, 2024, <https://arcprize.org/blog/oai-o3-pub-breakthrough>

³³LeGris, S., Vong, W. K., Lake, B. M., and Gureckis, T. M., 'H-ARC: A Robust Estimate of Human Performance on the Abstraction and Reasoning Corpus Benchmark', arXiv, 2024, <https://arxiv.org/pdf/2409.01374>

³⁴Ars Technica, 'OpenAI Announces o3 and o3-mini, Its Next Simulated Reasoning Models', Ars Technica, 2024, available at <https://arstechnica.com/information-technology/2024/12/openai-announces-o3-and-o3-mini-its-next-simulated-reasoning-models/>.

³⁵Goh, E. et al., 'Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial', JAMA Network Open, 2024, available at <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2825395>.

³⁶Hernström, V. et al., 'Screening Performance and Characteristics of Breast Cancer Detected in the Mammography Screening with Artificial Intelligence Trial (MASAI): A Randomised, Controlled, Parallel-Group, Non-Inferiority, Single-Blinded, Screening Accuracy Study', The Lancet, 2025, available at <https://www.thelancet.com/journals/landig/article/PIIS2589-7500%2824%2900267-X/fulltext>.

³⁷Dell'Acqua, F. et al., 'Navigating The Jagged Technological Frontier: Field Experimental Evidence Of The Effects Of AI On Knowledge Worker Productivity And Quality', Harvard Business School, 2023, available at <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>.

And AI's value is increasingly apparent across other professional and research domains. In contract review, [AI systems match or exceed expert lawyers](#)³⁸ in accuracy while dramatically reducing the time required. Meanwhile, a recent MIT study found that AI-assisted scientists "[discover 44% more materials, resulting in a 39% increase in patent filings and a 17% rise in downstream product innovation.](#)"³⁹ These tools have also [accelerated drug discovery](#)⁴⁰ and enabled unprecedented breakthroughs in protein-structure prediction, culminating in a [2024 Nobel Prize in Chemistry](#)⁴¹ for AlphaFold's achievements. Companies using AI tools report measurable gains, [averaging 25% productivity improvements across professional tasks](#)⁴².

Finally, [rapid uptake](#)⁴³ of AI tools by companies—despite known limitations—suggests they deliver real value beyond benchmark achievements.

In conclusion, while individual benchmarks are by no means perfect, we do not have to rely on them alone to assess AI capabilities. The combination of controlled human-AI comparisons (see Figure 2), successful real-world deployments, and measurable scientific and economic impact provides stronger evidence of genuine progress.

³⁸Choi, J.H., Monahan, A. and Schwarcz, D., 'Lawyering in the Age of Artificial Intelligence', SSRN, 2023, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4626276.

³⁹Toner-Rodgers, A., 'Artificial Intelligence, Scientific Discovery, and Product Innovation', NBER Conference Papers, 2024, available at https://conference.nber.org/conf_papers/f210475.pdf.

⁴⁰Jayatunga, M.K.P., Xie, W., Ruder, L., Schulze, U., and Meier, C., 'AI in small-molecule drug discovery: a coming wave?', Nature Reviews Drug Discovery, vol. 21, 2022, pp. 175–176, available at <https://www.nature.com/articles/d41573-022-00025-1>

⁴¹Financial Times, 'Google DeepMind Duo Share Nobel Chemistry Prize with US Biochemist', Financial Times, 2024, available at <https://www.ft.com/content/4541c07b-f5d8-40bd-b83c-12c0fd662bd9>.

⁴²Goldman Sachs, 'AI Is Showing 'Very Positive' Signs of Eventually Boosting GDP and Productivity', Goldman Sachs, 2024, available at <https://www.goldmansachs.com/insights/articles/AI-is-showing-very-positive-signs-of-boosting-gdp>.

⁴³McKinsey, 'The State of AI in 2024', McKinsey & Company, 2024, available at <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.



Figure 2: Comparison of AI vs human performance in multiple professional tasks; Source: CFG, using data from [medical diagnosis](#)⁴⁴, [consulting quality](#)⁴⁵ and [research output](#)⁴⁶ (in case of research output, numbers represent “discovery of new materials” and don’t have a maximum of 100% and thus the human baseline is set as 50%; in case of consulting quality, the indicator was human-rated score on a scale of 1–8 where the control group achieved score of 4.1)

⁴⁴Goh, E. et al., 'Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial', JAMA Network Open, 2024, available at <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2825395>.

⁴⁵Dell'Acqua, F. et al., 'Navigating The Jagged Technological Frontier: Field Experimental Evidence Of The Effects Of AI On Knowledge Worker Productivity And Quality', Harvard Business School, 2023, available at <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>.

⁴⁶Toner-Rodgers, A., 'Artificial Intelligence, Scientific Discovery, and Product Innovation', NBER Conference Papers, 2024, available at https://conference.nber.org/conf_papers/f210475.pdf.

AI still makes basic mistakes, surely that will limit its potential for transformation?

Critics often highlight AI's surprising errors in seemingly simple tasks, using these failures to question its reliability for critical applications. Yet despite these concerns, AI adoption is accelerating across industries, with businesses finding ways to extract value through human oversight and careful deployment strategies. Performance metrics show steady improvement in reliability, and the technology is proving useful even in high-stakes domains like healthcare, suggesting that while perfect reliability remains a challenge, it's not an insurmountable barrier to creating transformative impact.

For many everyday users, interactions with chatbots and open-access models reveal a striking discrepancy between the polished performance seen in controlled demos and the unpredictable behaviour in real-world use. Numerous social media posts, studies, and user reviews document instances where these systems [miscount letters in simple words](#)⁴⁷, [fail basic arithmetic](#)⁴⁸ or [make up facts](#)⁴⁹. This experience makes it hard to believe that AI models can indeed achieve the impressive impacts touted in headlines. While AI's capabilities are [expanding](#)⁵⁰, the technology has a [distinct set of strengths and weaknesses compared to humans](#)⁵¹ and its performance remains highly variable—excelling in tasks humans find difficult yet failing at some we find trivial. It's worth noting that this is an issue that far predates the current AI hype cycle - Moravec's Paradox, which states "it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility," was coined back in 1988⁵².

This fuels concerns about reliability: if AI stumbles on basic tasks, how can we trust it in critical applications? These mistakes aren't just embarrassing—they reveal a fundamental unpredictability that raises serious questions about AI deployment. [Some argue](#)⁵³ that until AI becomes consistently reliable, it won't see widespread adoption, especially in high-stakes fields like healthcare, infrastructure, or finance, where even rare mistakes can have serious consequences. If we can't trust AI to count letters correctly, how can we responsibly deploy it for complex, critical decisions?

⁴⁷OpenAI Developer Community, 'Incorrect Count of 'r' Characters in the Word 'strawberry' - Use Cases and Examples', OpenAI Developer Community, 2024, available at <https://community.openai.com/t/incorrect-count-of-r-characters-in-the-word-strawberry/829618>.

⁴⁸Reddit, 'ChatGPT getting basic math wrong', 2023, available at https://www.reddit.com/r/ChatGPT/comments/103bw12/chat_gpt_getting_basic_maths_wrong/.

⁴⁹Tuladhar, S.K., 'Two things ChatGPT can get wrong: math and facts. Why?', Medium, 2023, available at <https://medium.com/@k.sunman91/two-things-chatgpt-can-get-wrong-math-and-facts-why-b40e50dc68cd>

⁵⁰Google, 'Introducing Gemini 2.0 | Our Most Capable AI Model Yet', YouTube, 2024, available at <https://www.youtube.com/watch?v=Fs0t6SdODd8&t=45s>.

⁵¹Exponential View, 'From ChatGPT to a Billion Agents', Exponential View, 2024, available at <https://www.exponentialview.co/p/from-chatgpt-to-a-billion-agents>.

⁵²Moravec, H., *Mind Children*, Harvard University Press, 1988.

⁵³Marcus, G., 'What If Generative AI Turned Out to be a Dud?', Communications of the ACM, vol. 66, no. 3, 2023, available at <https://cacm.acm.org/blogcacm/what-if-generative-ai-turned-out-to-be-a-dud/>

While there is no disputing these errors persist, despite these concerns, AI adoption is actually accelerating. The latest [McKinsey survey](#)⁵⁴ linked in Figure 3 below shows 65% of respondents now regularly use generative AI—nearly double from ten months ago. Businesses are finding ways to extract value despite AI's imperfections, often by pairing it with human oversight or deploying it in lower-risk settings first. Yet a recent [Pew Research survey](#) indicates that the situation in the general public seems to be less enthusiastic— [only about a third of U.S. adults have ever used an AI chatbot](#)⁵⁵. This points to a more cautious or limited uptake among the broader public compared to reported within-organisations adoption figures, highlighting that while many organisations are forging ahead, mainstream adoption is still emerging, varying considerably by context and user needs. Nonetheless, if these reported levels of AI adoption within businesses prove to be broadly accurate, it suggests that, even without perfect reliability in high-stakes settings, AI's use in low-stakes contexts could still be transformative for society and generate substantial economic value.

Interestingly, AI is proving to be useful even in high-stakes domains. In medical diagnosis, AI models already perform at a higher [accuracy than human doctors](#)⁵⁶, and another study found [ChatGPT to be nearly 72% accurate across all medical specialties](#)⁵⁷ and phases of clinical care including generating possible diagnoses and care management decisions, and 77% accurate in making final diagnoses in an experiment comparing doctors with and without AI assistance. Several specialised AI-enabled medical tools [were also approved for use by the US Food and Drug Administration](#)⁵⁸, suggesting the reliability was high enough for the given use cases.

Moreover, reliability metrics show steady improvement. One key aspect of AI reliability is how often models make up information that isn't true—a phenomenon known as "hallucination". On [Vectara's Hughes Hallucination Leaderboard](#)⁵⁹, where AI models are given a simple goal of summarising a provided document, OpenAI's GPT-3.5 had a hallucination rate of 3.5% in November 2023, but as of January 2025, the firm's later model GPT-4 scored 1.8% and latest o3-mini-high-reasoning model 0.8%. However, broader evaluations in more open-ended scenarios (i.e. not based on a specific document that the AI model should refer to) do not always show a clear downward trend. OpenAI reports that while o1 outperformed GPT-4 on its internal hallucination benchmarks, testers observed

⁵⁴McKinsey, 'The State of AI in 2024', McKinsey & Company, 2024, available at <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

⁵⁵McClain, C., Kennedy, B., Gottfried, J., Anderson, M., and Pasquini, G., 'How the U.S. Public and AI Experts View Artificial Intelligence', Pew Research Center, 2025, available at <https://www.pewresearch.org/2025/04/03/artificial-intelligence-in-daily-life-views-and-experiences/>

⁵⁶Goh, E. et al., 'Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial', JAMA Network Open, 2024, available at <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2825395>.

⁵⁷Rao, A. et al., 'Assessing the Utility of ChatGPT Throughout the Entire Clinical Workflow: Development and Usability Study', Journal of Medical Internet Research, 2023, available at <https://www.jmir.org/2023/1/e48659/>.

⁵⁸U.S. Food and Drug Administration (FDA), 'Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices', FDA, 2024, available at <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-ai-ml-enabled-medical-devices>.

⁵⁹Vectara, 'Hallucination Leaderboard', GitHub, n.d., <https://github.com/vectara/hallucination-leaderboard>.

that it [sometimes generated more hallucinations](#)⁶⁰—particularly elaborate but incorrect responses that appeared more convincing. On TruthfulQA, a benchmark measuring model truthfulness, the performance of the best AI model has improved [from 58% accuracy in 2022](#)⁶¹ when the benchmark was published to [91.1% as of January 2025](#)⁶², which is just short of human performance at 94%.

In conclusion, while AI's lack of reliability remains a challenge—especially in high-stakes settings—real-world trends suggest they are not an insurmountable barrier to widespread adoption and creation of impact. Further, some paths to transformative impact—like accelerating scientific progress—don't require widespread adoption or a balanced human-like skillset and might not require perfect reliability either.




Figure 3: Percentage of respondents to the 2024 McKinsey Global Survey from organisations that have adopted AI in at least 1 business function; showing that generative AI adoption worldwide has increased dramatically from 2023–24; Source: Adapted from [McKinsey Global Survey 2024](#)⁶³

⁶⁰Jones, N., "'In awe'": scientists impressed by latest ChatGPT model o1', Nature, 2024, available at <https://www.nature.com/articles/d41586-024-03169-9>.

⁶¹Lin, S., Hilton, J., and Evans, O., 'TruthfulQA: Measuring How Models Mimic Human Falsehoods', arXiv, 2021, <https://arxiv.org/abs/2109.07958>.

⁶²Evans, O., Chua, J., and Lin, S., 'New, improved multiple-choice TruthfulQA', LessWrong, 2025, available at <https://www.lesswrong.com/posts/Bunfwz6JsNd44kgLT/new-improved-multiple-choice-truthfulqa>.

⁶³McKinsey, 'The State of AI in 2024', McKinsey & Company, 2024, available at <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

We've seen tech hype cycles before—why is this time different?

Sceptics dismiss current AI developments by drawing parallels to past innovations—from non-fungible tokens (NFTs) to virtual reality and the metaverse—that generated enormous excitement before falling short of transformative predictions. However, AI fundamentally differs from previous technologies by enhancing cognitive capabilities rather than just physical ones or information access. Its distinctive characteristics include the ability to improve itself through automated research, unprecedented adaptability across diverse tasks through fine-tuning and prompting, and unique deployment advantages in our digital age that allow it to reach millions of users instantly through existing networks and infrastructure without requiring extensive physical installations or hardware upgrades. The combination of these features suggests a qualitatively different kind of technological revolution with potential for broader and faster impact than historical patterns would indicate.

From virtual reality to the blockchain to the metaverse, recent history is littered with technologies that generated enormous excitement and investment [before falling short of transformative predictions](#)⁶⁴. Some observers argue that AI will follow a similar pattern—generating massive hype and investment before reality sets in and more modest capabilities and impact are revealed. While this scepticism draws on valid historical examples, it may overlook fundamental differences in AI's nature and potential.

Throughout history, major technological revolutions have typically [augmented our physical abilities](#) (like steam engine and electricity powering machines to multiply human strength)⁶⁵ or [improved information access and exchange](#) (like the internet and smartphones connecting people to vast knowledge networks)⁶⁶. While humans have created impactful tools for the cognitive domain as well, they primarily functioned as extensions of our own thinking—calculators execute operations we defined, and writing preserves thoughts we had already formed. What distinguishes modern AI is its ability to independently generate cognitive outputs that can appear new or unexpected—while some argue [it's merely recombining existing information](#)⁶⁷, others see emergent patterns and insights that [weren't explicitly programmed](#)⁶⁸. Unlike electricity, which powers devices but doesn't decide how to use that power, or the internet, which connects information but doesn't generate new knowledge on its own, AI systems rely on human input yet can produce reasoning chains that we did not directly design or anticipate.

⁶⁴Floridi, L., 'Why the AI Hype is Another Tech Bubble', *Philosophy & Technology*, Springer, 2024, available at <https://link.springer.com/article/10.1007/s13347-024-00817-w>.

⁶⁵ National History Education Clearinghouse, 'Innovation and Technology in the 19th Century', Teachinghistory.org, n.d., available at <https://teachinghistory.org/history-content/ask-a-historian/24470>.

⁶⁶ National Research Council, Global Affairs, and Panel on the Impact of Information Technology on the Future of the Research University, *Preparing for the Revolution: Information Technology and the Future of the Research University*, Washington, DC, National Academies Press, 2002, available at <https://nap.nationalacademies.org/read/10545/chapter/3#7>.

⁶⁷ Bender, E.M. et al., 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?', in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 610–623, available at

https://dl.acm.org/doi/pdf/10.1145/3442188.3445922?utm_source=miragenews&utm_medium=miragenews&utm_campaign=news.

⁶⁸ Wei, J. et al., 'Emergent Abilities of Large Language Models', arXiv, 2022, available at <https://arxiv.org/abs/2206.07682>.

Whether it's truly understanding and creating meaning or merely recombining patterns from its training data, this shift—away from purely speeding up calculations toward contributing to the cognitive process itself—suggests a more active role for technology in shaping ideas. Although not entirely autonomous, the ability of AI to engage in what seems like genuine problem-solving rather than passively amplifying physical effort or communication channels marks a profound shift in how technology might impact human civilisation.

Traditional technologies also don't accelerate their own development; they rely on human engineers to design upgrades. AI, by contrast, can be used to develop better AI systems by improving performance of AI chips, creating more synthetic data for AI training, and coming up with more efficient algorithms to learn from the data. This creates feedback loops that could accelerate progress (see Figure 4). While AI is not yet able to autonomously design more capable AI systems, it is already speeding up the pace of development. For example, [Google reports](#)⁶⁹ that AI now assists in writing over 25% of their code, which could include work that advances AI itself. [A recent study](#)⁷⁰ has found that today's frontier AI models can already speed up typical AI R&D tasks by almost 2x and are [outperforming human researchers](#)⁷¹ in AI development tasks taking fewer than two hours. This self-perpetuating feedback loop has few parallels in past technological revolutions, where progress typically came in slower increments.

AI represents a new level of technological flexibility and adaptability. While past breakthroughs often excelled in one domain, AI can be rapidly adapted to new tasks through fine-tuning or even mere prompting, without requiring fundamental redesign. The same AI model can jump from writing legal briefs to optimising research pipelines to generating 3D images. This unprecedented adaptability means that innovations in AI capabilities can quickly propagate across multiple sectors and applications, amplifying their impact.

The economics of AI deployment are uniquely suited to our digital age. We live in a world where much of human activity—from commerce to communication, from work to entertainment—already takes place through digital interfaces. This existing digital infrastructure allows AI to integrate directly into these workflows and activities, amplifying its impact on both the digital and physical worlds. Combined with [sharply decreasing costs of usage](#)⁷², this enables unprecedented speed of adoption. ChatGPT, for instance, became [the fastest-growing consumer application in history](#)⁷³, reaching 100 million users in just two months.

⁶⁹Google, 'Alphabet Q3 Earnings Call: CEO Sundar Pichai's Remarks', Google (blog), 2024, <https://blog.google/inside-google/message-ceo/alphabet-earnings-q3-2024/#full-stack-approach>.

⁷⁰Lin, T., 'Measuring Automated Kernel Engineering', METR, 2025, <https://metr.org/blog/2025-02-14-measuring-automated-kernel-engineering/>.

⁷¹Wijk, H. et al., 'Re-bench: Evaluating Frontier AI R&D Capabilities Of Language Model Agents Against Human Experts', METR, 2024, available at <https://metr.org/blog/2024-11-22-evaluating-r-d-capabilities-of-llms/>.

⁷²Tsarathustra, 'Dane Vahey of OpenAI Says the Cost...', X, 2024, available at <https://x.com/tsarnick/status/1838401544557072751>.

⁷³Hu, K., 'ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note', Reuters, 2023, available at <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>.

In conclusion, while some individual [companies](#)⁷⁴ or [applications](#)⁷⁵ may indeed be overhyped, dismissing current AI developments as "just another fad" overlooks crucial differences in the technology's nature, deployment environment and potential for scale of the impact.

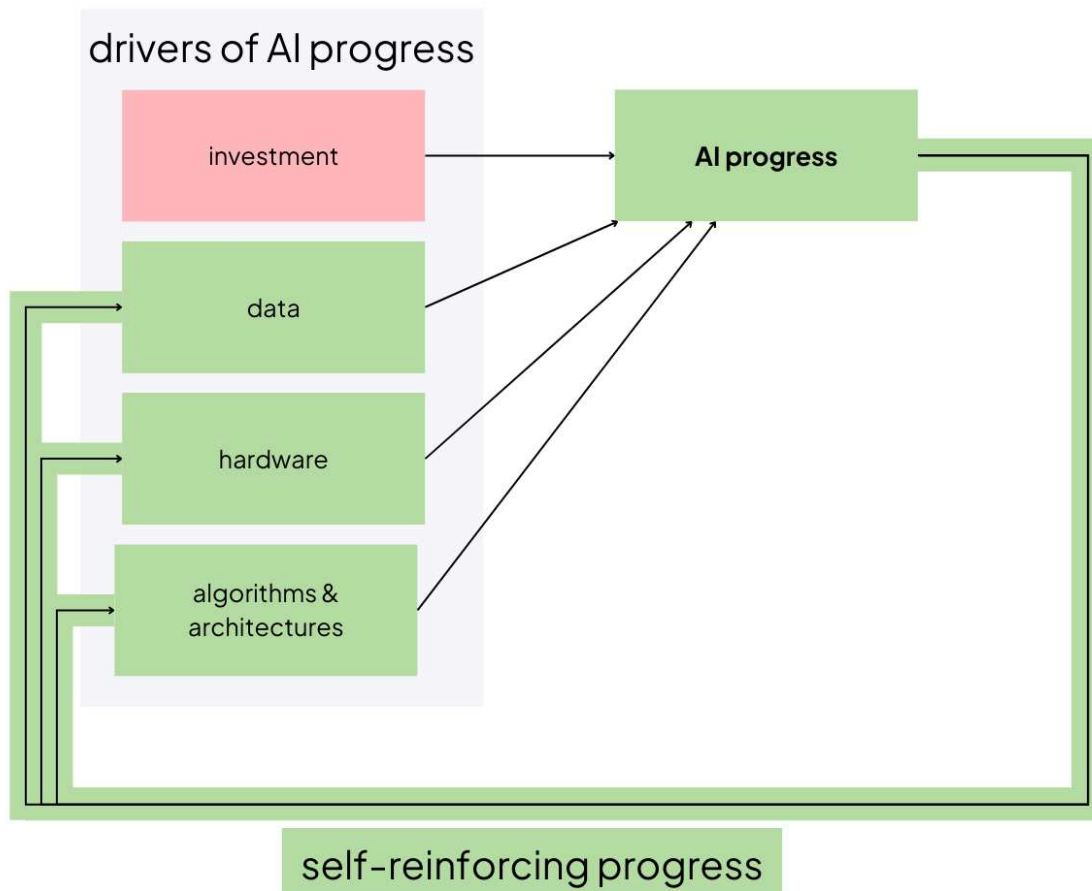


Figure 4: Conceptual diagram of how AI progress can fuel further AI progress. For three of the four main drivers—data, hardware, and algorithms & architectures—each advancement in AI makes these drivers stronger, which in turn accelerates future AI breakthroughs. Source: CFG

⁷⁴Olson, P., 'Nearly Half Of All 'AI Startups' Are Cashing In On Hype', Forbes, 4 March 2019, available at <https://www.forbes.com/sites/parmyolson/2019/03/04/nearly-half-of-all-ai-startups-are-cashing-in-on-hype/>.

⁷⁵Benaiche, E. (2023) 'AI Washing', Goaland, available at <https://www.goaland.com/blog/AI/ai-washing>.

Is AI hitting fundamental limits or will new methods unlock the next leap?

Sceptics point to potential resource constraints in AI development, from high-quality data availability to computing power and electricity supply needs. A deeper concern, raised by prominent experts like Yann LeCun and François Chollet, suggests that the fundamental architecture of large language models may be approaching its capability ceiling, with pre-training gains showing signs of diminishing returns. However, the field demonstrates remarkable adaptability: as traditional pre-training approaches a plateau, new methodologies like post-training optimization and runtime improvements are already showing promising results, with models achieving unprecedented performance on complex reasoning tasks. While future transitions between approaches might not be equally smooth, the field's proven ability to pivot and innovate suggests we're entering a new phase of development rather than hitting absolute limits.

When [critics](#)⁷⁶ [argue](#)⁷⁷ [that AI is hitting](#)⁷⁸ [fundamental limits](#)⁷⁹ they're making a nuanced technical argument about the nature of current AI progress. The concern isn't simply that progress will slow down—it's that we're approaching inherent limitations of the dominant approach to AI development.

One concern is whether the massive resources that have fuelled AI's progress—particularly data, computing power, and electricity supply—will remain abundant. Historically, most AI capabilities came from [scaling](#)⁸⁰ up training efforts with more data and more chips, requiring more power. This approach creates potential bottlenecks, from limited training data and electricity supply to specialised chip production constraints. While it's true that training frontier models demands tremendous resources, [recent research](#)⁸¹ suggests all these necessary inputs are likely to be available at least until 2030 at the current rate of scaling.

Some experts raise an even deeper concern: the architecture of large language models (LLM)—the most successful architecture of AI development in recent years—may be reaching its limits. [They mention](#)⁸² various reasons for why the LLM architecture might be limited: reliance on statistical

⁷⁶Marcus, G., 'Deep Learning Is Hitting a Wall', Nautilus, 2024, available at <https://nautil.us/deep-learning-is-hitting-a-wall-238440/>.

⁷⁷Patel, D., 'Francois Chollet - Why The Biggest AI Models Can't Solve Simple Puzzles', YouTube, 2024, available at <https://www.youtube.com/watch?v=UakqL6Pj9xo>.

⁷⁸Piquard, A., 'Chatbots Are Like Parrots, They Repeat Without Understanding', Le Monde, 2024, available at https://www.lemonde.fr/en/economy/article/2024/10/07/chatbots-are-like-parrots-they-repeat-without-understanding_6728523_19.html.

⁷⁹LeCun, Y., 'About the raging debate regarding the future of AI', Facebook post, 2024, available at <https://www.facebook.com/yann.lecun/posts/10158256523332143>.

⁸⁰Hestness, J. et al. (2017) 'Deep Learning Scaling Is Predictable, Empirically', ArXiv, available at <https://arxiv.org/abs/1712.00409>.

⁸¹Sevilla, J. et al., 'Can AI Scaling Continue Through 2030?', Epoch AI, 2024, available at <https://epoch.ai/blog/can-ai-scaling-continue-through-2030>.

⁸²Ball, D.W., 'The Era of Predictive AI Is Almost Over', The New Atlantis, no. 77, 2024, pp. 26–35, available at <https://www.thenewatlantis.com/publications/the-era-of-predictive-ai-is-almost-over>.

correlations rather than true causal understanding; lack of a model of the world that gives meaning to the text; lack of ability to improvise and adapt to novel, unforeseen situations as some of the examples. Recent [indications](#)⁸³ suggest that capability gains from pre-training may be [slowing](#)⁸⁴, prompting some to speculate that deep learning might be [hitting a wall](#)⁸⁵. Under this view, more scale alone would yield only marginal improvements, and fully new architectures—or entire paradigms—might be needed to push AI's capabilities significantly further.

Yet new training methods are already stepping in where pre-training seems to be losing steam. [“Post-training” and runtime optimisations](#)⁸⁶—enabling models to deliberate longer before producing output—are showing [remarkable promise](#)⁸⁷. OpenAI's latest models which leveraged this approach jumped [from 2% to 25%](#)⁸⁸ accuracy on the FrontierMath benchmark⁸⁹ and [surpassed human performance](#)⁹⁰ on ARC-AGI with scores of 75.7%. Meanwhile, the first open-source model using the same architecture, [DeepSeek R1](#)⁹¹, has garnered significant attention for reducing costs while boosting capabilities.

As Figure 5 illustrates, technological progress often follows a pattern of successive S-curves, where each approach eventually reaches diminishing returns only to be replaced by a new paradigm that overcomes previous limitations. What critics identify as "hitting a fundamental limit" could simply be the flattening top of one curve before the next approach accelerates progress again. We don't know

⁸³Hu, K. and Tong, A., 'OpenAI and others seek new path to smarter AI as current methods hit limitations', Reuters, 2024, available at <https://www.reuters.com/technology/artificial-intelligence/openai-rivals-seek-new-path-smarter-ai-current-methods-hit-limitations-2024-11-11/>.

⁸⁴Metz, R. et al., 'OpenAI, Google and Anthropic Are Struggling to Build More Advanced AI', Bloomberg, 2024, available at <https://www.bloomberg.com/news/articles/2024-11-13/openai-google-and-anthropic-are-struggling-to-build-more-advanced-ai>.

⁸⁵Marcus, G., 'Folks, Game Over....', X, 2024, available at <https://x.com/GaryMarcus/status/1855382564015689959>.

⁸⁶Juijn, D., 'No, Deep Learning Is (probably) Not Hitting A Wall', Centre for Future Generations, 2024, available at <https://cfg.eu/no-deep-learning-is-probably-not-hitting-a-wall/>.

⁸⁷Janků, D., Reddel, F., Graabak, J. and Reddel, M., 'Beyond the AI Hype – A Critical Assessment of AI's Transformative Potential', Centre for Future Generations, 2025, available at: <https://cfg.eu/beyond-the-ai-hype>.

⁸⁸Ars Technica, 'OpenAI Announces o3 and o3-mini, Its Next Simulated Reasoning Models', Ars Technica, 2024, available at <https://arstechnica.com/information-technology/2024/12/openai-announces-o3-and-o3-mini-its-next-simulated-reasoning-models/>.

⁸⁹Note that there is a recent controversy regarding OpenAI financing creation of this benchmark and having access to some of the data, thus potentially being able to optimise their models for high performance on this benchmark, making their impressive achievements invalid. The investigation whether they did use the benchmark data in this way is currently ongoing. More context can be found at: <https://www.searchenginejournal.com/openai-secretly-funded-frontiermath-benchmarking-dataset/537760>.

⁹⁰'OpenAI o3 Breakthrough High Score On ARC-AGI-pub', ARC Prize, 2024, <https://arcprize.org/blog/oai-o3-pub-breakthrough>.

⁹¹DeepSeek, 'DeepSeek-R1 Release', DeepSeek API Documentation, 20 January 2025, available at <https://api-docs.deepseek.com/news/news250120>.

the ultimate scaling potential of these new approaches; they, too, may encounter their own limits. The main question isn't whether these methods will hit limits, but how quickly new paradigms emerge once they do. The AI field has already demonstrated its ability to pivot rapidly when pre-training showed diminishing returns. It remains to be seen how smooth future transitions between approaches will be.

In conclusion, while concerns about current AI paradigms—particularly those reliant on large language models—hitting a wall deserve serious consideration, current evidence suggests we are entering a new phase of development rather than approaching absolute barriers. New training methodologies and optimisation techniques are showing tangible results, with significant performance improvements in areas where traditional pre-training appeared to plateau. However, important questions remain about the long-term scalability of these new approaches, the ultimate limitations of current architectures, and whether future transitions between paradigms will be as smooth as recent ones. While some methods may indeed be reaching their limits, the field's proven adaptability and the emergence of promising new techniques suggest that claims of fundamental barriers to progress may be premature.

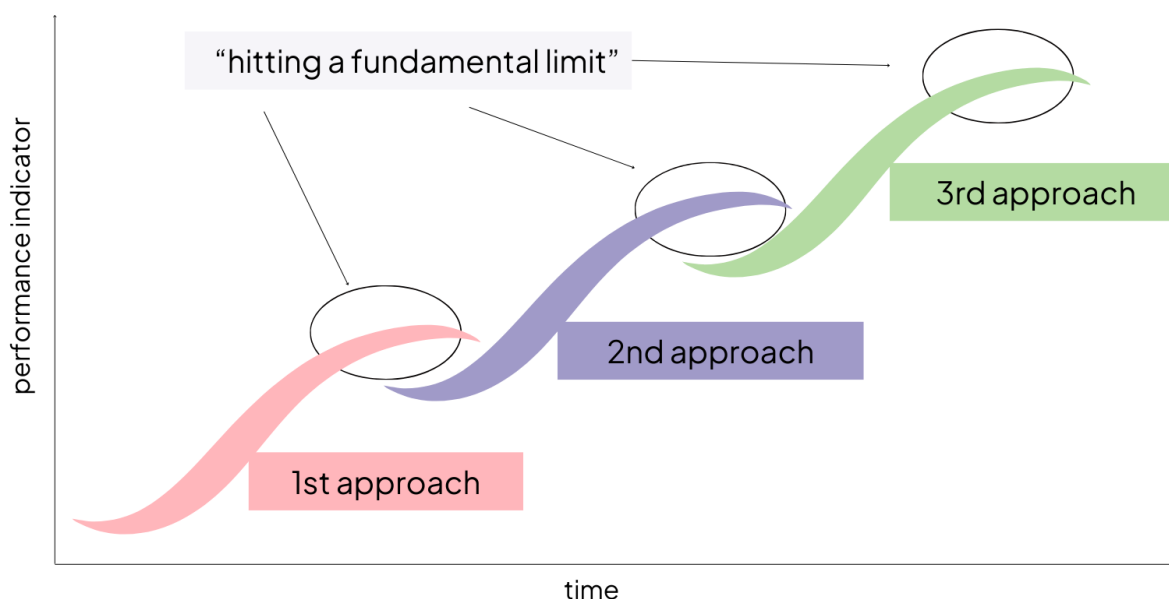


Figure 5: Successive S-curves in technological development; showing the conceptual idea of how new paradigms might offset the tapering gains of previous paradigms, sustaining overall progress. Source: CFG, inspired by Priestley, 2020⁹²

⁹²Priestley, M., Sluckin, T.J. and Tiropanis, T., 'Innovation on the Web: the end of the S-curve?', Internet Histories, vol. 4, 2020, available at https://www.researchgate.net/publication/340525808_Innovation_on_the_Web_the_end_of_the_S-curve.

Is transformative AI impact decades away or is change coming much faster than expected?

Many argue that meaningful transformation requires more than technical breakthroughs, pointing to the historical pattern of revolutionary technologies taking decades to realize their potential. This is broadly true but several factors distinguish AI's trajectory: it can assist in its own development, creating potential feedback loops for accelerated progress; it can leverage existing digital infrastructure for rapid deployment, avoiding the need for extensive new physical installations; and it's attracting unprecedented levels of investment and regulatory attention. Expert surveys show a dramatic shortening of predictions of when transformative AI arrives, with many researchers and industry leaders updating their estimates from several decades to a few years. Combined with measurable gains in real-world applications and rapidly growing adoption rates, these factors suggest we should seriously consider the possibility of transformative impact arriving much sooner than historical patterns would indicate.

Even if today's obstacles aren't fundamental, many still expect a slow journey to truly transformative impact⁹³. This viewpoint is reinforced by historical precedents—many revolutionary technologies took decades to realise their potential, and some believe AI will be no different. The thinking is that real impact requires both additional technical and design breakthroughs that get the product from the lab prototype to a tool useful in the real world, as well as large-scale adoption that is a result of many factors, including regulation, industry standards, workforce training, infrastructure upgrades, and countless other intersecting elements. In many cases, it's the alignment of these moving parts—[not just technical progress](#)—that determines how fast a technology spreads, as explored in recent work on how diffusion drives economic competition and national advantage⁹⁴.

On the technical side, progress suggests we might actually be underestimating AI's speed. In just two years, image-generation tools have progressed from blurry, low-detail outputs to photorealistic quality (see Figure 6), while [similarly steep gains](#)⁹⁵ appear in handwriting recognition, competition-level mathematics, and [other tasks](#) (see Figure 8)⁹⁶. Such rapid, broad improvements break with historical patterns, pointing to a much faster overall pace of AI advancement.

⁹³ We define transformative impact as “radical changes in welfare, wealth or power”, in line with Dafoe, A., 'AI Governance: A Research Agenda', Centre for the Governance of AI, Future of Humanity Institute, University of Oxford, 2018, available at <https://cdn.governance.ai/GovAI-Research-Agenda.pdf>.

⁹⁴ Ding, J., Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition, Princeton, NJ, Princeton University Press, 2024, available at <https://press.princeton.edu/books/paperback/9780691260341/technology-and-the-rise-of-great-powers?srsId=AfmBOopELBKgsfq0tTJYOzxeFzxfatALnpKgDfT-I5q4MIIIViWLnGjl>.

⁹⁵ UK Government, 'International AI Safety Report', UK Government, 2025, https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.

⁹⁶ Janků, D., Reddel, F., Graabak, J. and Reddel, M., 'Beyond the AI Hype – A Critical Assessment of AI's Transformative Potential', Centre for Future Generations, 2025, available at: <https://cfg.eu/beyond-the-ai-hype>. Section on Capabilities.



Figure 6: Example of progress in Midjourney’s image generation capabilities over the timespan of 2 years; showing how fast capabilities in AI image generation have progressed. Source: adapted from [Stanford’s 2024 AI Index Report, p.96](#)⁹⁷

Moreover, AI’s potential for self-improvement accelerates development speed further. Unlike earlier technologies, AI can automate parts of its own development, creating a feedback loop that, where AI helps develop better AI systems, potentially leading to much faster progress than historical patterns would suggest. AI research automation is also special compared to other industries in its feasibility—in large part, it relies on coding and software development, capabilities in which current [AI systems are already strongest](#)⁹⁸. Indeed, current AI systems are [already outperforming human researchers](#)⁹⁹ in tasks taking less than 2 hours and many people predict that [some form of full AI research automation](#)¹⁰⁰ might appear in 2028.

On the adoption side, the speed might also be faster compared to historical precedents. Patterns in broader technological adoption show that modern technologies are much quicker to be adopted (see Figure 7), perhaps partly thanks to the phenomenon of [technological convergence](#)¹⁰¹ when originally unrelated technologies become more closely integrated and even unified as they develop and advance.

⁹⁷UK Government, 'International AI Safety Report', UK Government, 2025, https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.

⁹⁸Krishnan, N., 'OpenAI's O3: A New Frontier in AI Reasoning Models', Towards AI, 2024, <https://pub.towardsai.net/openais-o3-a-new-frontier-in-ai-reasoning-models-a02999246ffe>.

⁹⁹Wijk, H. et al., 'Re-bench: Evaluating Frontier AI R&D Capabilities Of Language Model Agents Against Human Experts', METR, 2024, available at <https://metr.org/blog/2024-11-22-evaluating-r-d-capabilities-of-llms/>.

¹⁰⁰Metaculus, 'When will AIs program programs that can program AIs?', Metaculus, n.d., <https://www.metaculus.com/questions/406/date-ais-capable-of-developing-ai-software/>.

¹⁰¹Bainbridge, W.S. and Roco, M.C., 'Science and technology convergence: with emphasis for nanotechnology-inspired convergence', Journal of Nanoparticle Research, vol. 18, no. 7, 2016, article no. 211, available at <https://link.springer.com/article/10.1007/s11051-016-3520-0>.

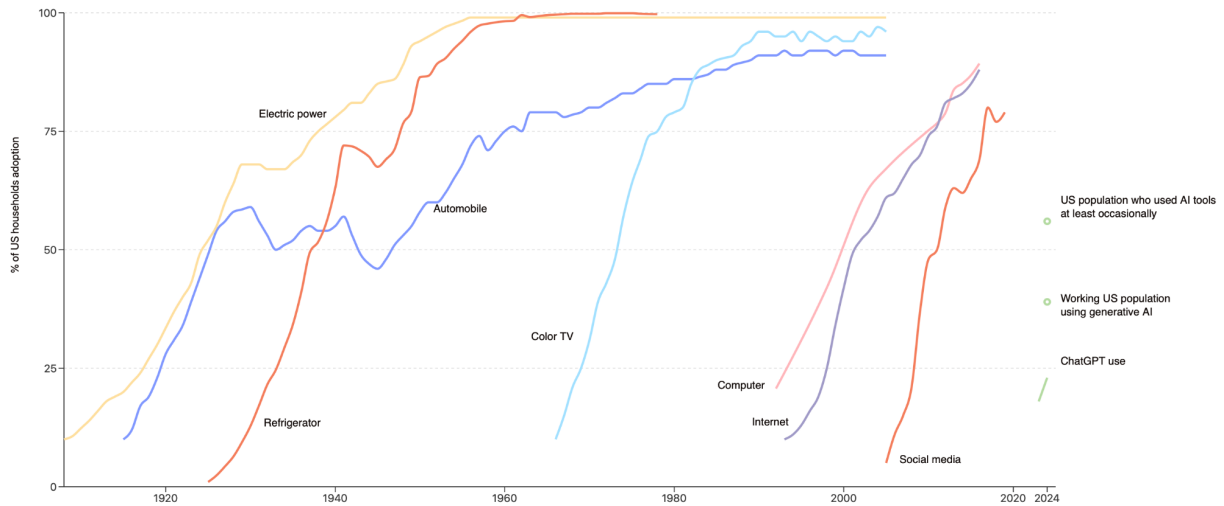


Figure 7: Adoption rates of various technologies and modern AI; showing that modern technologies are quicker to be adopted and AI adoption specifically is unusually fast. Source: [Our World in Data](#)¹⁰² for all technologies aside from AI, [Pew Research](#)¹⁰³ for ChatGPT use, [Bick et al 2024](#)¹⁰⁴ for Working US population and [YouGov](#)¹⁰⁵ for US population using AI occasionally

Furthermore, the AI sector has an advantage in that it can deliver impact via existing digital infrastructure. While AI requires substantial investment in data centres and computational resources, the crucial distinction is that these investments happen primarily on the provider side¹⁰⁶. From the user perspective, AI capabilities can be accessed through ordinary computers, smartphones, and internet connections already in place. Unlike earlier transformative technologies—electricity required new wiring in every building, the internet demanded new cables and modems in homes—AI services are immediately accessible to anyone with existing digital devices. This accessibility pattern explains why AI adoption can spread rapidly across diverse sectors.

¹⁰²Our World in Data, 'Share of United States Households Using Specific Technologies', Our World in Data, n.d., available at <https://ourworldindata.org/grapher/technology-adoption-by-households-in-the-united-states>.

¹⁰³Pew Research Center, 'Americans Increasingly Using ChatGPT, but Few Trust Its 2024 Election Information', Pew Research Center, 2024, available at <https://www.pewresearch.org/short-reads/2024/03/26/americans-use-of-chatgpt-is-ticking-up-but-few-trust-its-election-information/>.

¹⁰⁴Bick, A., Blandin, A., and Deming, D.J., 'The Rapid Adoption of Generative AI', 2024, available at https://889099f7-c025-4d8a-9e78-9d2a22e8040f.usrfiles.com/ugd/889099_fa955a5012dd45abb84dd647ce02de95.pdf.

¹⁰⁵YouGov, 'Do Americans Think AI Will Have a Positive or Negative Impact on Society?', YouGov, 2025, available at <https://today.yougov.com/technology/articles/51368-do-americans-think-ai-will-have-positive-or-negative-impact-society-artificial-intelligence-poll>.

¹⁰⁶ Though building those remote datacenters could be challenging and costly undertaking, which has already taken a toll on energy grids (see Baraniuk, C., 'Electricity grids creak as AI demands soar', BBC News, 2024, available at <https://www.bbc.com/news/articles/cj5ll89dy2mo>.)

This digital-first flexibility is already reflected in AI's rapid uptake: ChatGPT [became the fastest-growing app in history](#)¹⁰⁷, reaching millions of users in days instead of years, now totalling [400 million weekly users](#)¹⁰⁸. While AI could transform all knowledge work rapidly, transforming the physical tasks might take longer due to additional challenges in mechanical engineering and real-world adaptability. Yet even in its primarily cognitive applications, AI has the potential to influence our physical world by helping optimise industrial processes, assisting with human labor coordination, and improving existing machinery control systems.

Meanwhile, providers are leaving nothing to chance: big tech companies are investing heavily in new infrastructure dedicated to AI training and deployment, further expanding adoption. Current projections put these investments at around [\\$240 billion in 2024 alone](#)¹⁰⁹.

Across academia, industry, and consulting, expert predictions for when transformative AI arrives have shifted dramatically closer. A [2023 survey](#)¹¹⁰ of 2,778 researchers who had published in top-tier AI journals showed that their predictions for when AI reaches transformative milestones have come significantly closer recently, perhaps informed by the high speed of AI development. This represents the largest update in timeline predictions in the history of such surveys, suggesting that even experts are repeatedly surprised by the pace of progress. Industry leaders closest to AI development—though potentially motivated to emphasise its promise—are becoming particularly bullish. [Many of them](#)¹¹¹ [expect](#)¹¹² AI system that performs at or above human-level across all cognitive tasks to be developed in the next 5 years, and some notable significant impact materializing [as soon as in 2025](#)¹¹³.

¹⁰⁷Hu, K., 'ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note', Reuters, 2023, available at <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>.

¹⁰⁸Reuters, '400 Million Weekly Users', Reuters, 2025, available at <https://www.reuters.com/technology/artificial-intelligence/openai-weekly-active-users-surpass-400-million-2025-02-20/>.

¹⁰⁹Forbes, 'AI Spending to Exceed a Quarter Trillion Next Year', Forbes, 2024, available at <https://www.forbes.com/sites/bethkindig/2024/11/14/ai-spending-to-exceed-a-quarter-trillion-next-year/>.

¹¹⁰Grace, K. et al., 'Thousands of AI Authors on the Future of AI', arXiv, 2024, <https://arxiv.org/abs/2401.02843>.

¹¹¹Varanasi, L., 'Here's How Far We Are From AGI, According to the People Developing It', Business Insider, 2024, available at <https://www.businessinsider.com/agi-predictions-sam-altman-dario-amodei-geoffrey-hinton-demis-hassabis-2024-11>.

¹¹²Fridman, L., 'Dario Amodei: Anthropic CEO On Claude, AGI & The Future Of AI & Humanity | Lex Fridman Podcast #452', YouTube, 2024, <https://www.youtube.com/watch?v=ugvHCXC0mm4>.

¹¹³OpenAI's CEO Sam Altman, for example, has said that in 2025, we may see the first AI agents "join the workforce" and materially boost company output, adding that ["AGI will probably get developed during \[Trump's\] term."](#); Source: Pillay, T., 'How OpenAI's Sam Altman Is Thinking About AGI and Superintelligence in 2025', TIME, 2025, <https://time.com/7205596/sam-altman-superintelligence-agi/>.

This aligns with predictions of more independent experts as well, such as Turing and Nobel Award Laureate and “Godfather of AI” Geoffrey Hinton [saying](#)¹¹⁴ that AI could be as little as five years away from surpassing human intelligence, a significant switch from his previous predictions of 30–50 years. Even McKinsey, known for its pragmatic market analysis, [highlights that](#)¹¹⁵ “increasing investments in these tools could result in agentic systems achieving notable milestones and being deployed at scale over the next few years”.

In conclusion, even if earlier technologies took decades to become transformative, AI’s unique characteristics, together with the established technological ecosystem it inhabits, suggest much faster progress. Concerns about resource constraints—from data availability to computing power—appear manageable [at least through 2030 based on current projections](#)¹¹⁶, and governments have [accelerated policy responses](#)¹¹⁷—nearly doubling new regulatory initiatives year over year. Investment is also at unprecedented scales, with projects like the US government-supported [\\$500 billion Stargate plan](#)¹¹⁸, [China’s one trillion Yuan](#)¹¹⁹ (approximately \$138 billion) “AI Industry Development Action Plan” and [EU with a total of over €200 billion](#)¹²⁰, demonstrating serious commitment to building the necessary infrastructure. Meanwhile, AI’s self-improving nature, digital-first deployment, and unprecedented funding levels create the potential for much faster transformation than historical patterns would suggest. This acceleration in both capabilities and adoption means we should seriously consider the possibility of transformative impact within years rather than decades—and prepare accordingly.

¹¹⁴Hinton, G. [@geoffreyhinton], Twitter post, 2023, available at <https://x.com/geoffreyhinton/status/1653687894534504451>.

¹¹⁵Yee, L., Chui, M., and Roberts, R., with Xu, S., ‘Why agents are the next frontier of generative AI’, McKinsey Digital, 2024, <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai>.

¹¹⁶Sevilla, J. et al., ‘Can AI Scaling Continue Through 2030?’, Epoch AI, 2024, available at <https://epoch.ai/blog/can-ai-scaling-continue-through-2030>.

¹¹⁷Stanford HAI, ‘HAI AI-Index-Report-2024’, Stanford HAI, 2024, available at https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf.

¹¹⁸Silva, J., Sherman, N., and Rahman-Jones, I., ‘Stargate: Tech giants announce AI plan worth up to \$500bn’, BBC News, 2024, available at <https://www.bbc.com/news/articles/cy4m84d2xz2o>.

¹¹⁹Parasnis, S., ‘Bank of China announces 1 Trillion Yuan AI Industry investment’, Medianama, 2025, available at <https://www.medianama.com/2025/01/223-bank-of-china-announces-1-trillion-yuan-ai-industry-investment/>.

¹²⁰Wold, J.W., ‘Von der Leyen Launches World’s Largest Public-Private Partnership To Win AI Race’, EURACTIV, 2025, available at <https://www.euractiv.com/section/tech/news/von-der-leyen-launches-worlds-largest-public-private-partnership-to-win-ai-race/>.

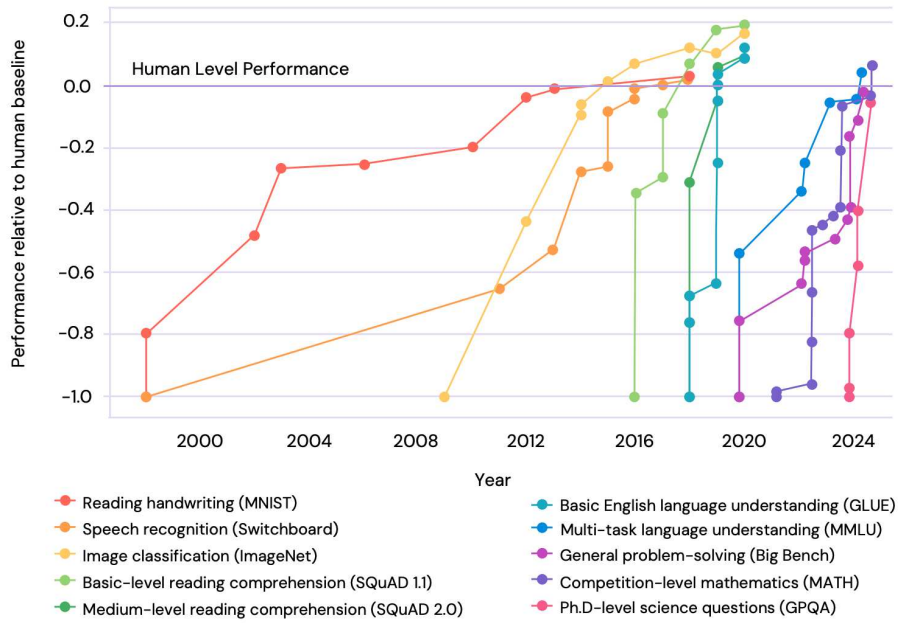


Figure 8: AI performance on various skills over time compared to human performance; showing unprecedented and accelerating speed of AI capability development. Source: [International AI safety report \(2025, p. 49\)](https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf)¹²¹

¹²¹UK Government, 'International AI Safety Report', UK Government, 2025, https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.

Acknowledgements

We would like to express our sincere gratitude to Aaron Maniam, David Timis, Kristina Fort, Leonardo Quattrucci, Pawel Swieboda, and many of our CFG colleagues for their valuable feedback and insights on this report. Their thoughtful comments and suggestions have significantly improved the quality of our work. Any remaining errors or omissions are our own.

Contact us

We welcome engagement from policymakers, researchers, industry experts, and anyone interested in separating AI reality from the hype. Whether you have questions, feedback, or would like to discuss our findings further, we'd love to hear from you.

 **Email:** d.janku@cfg.eu

 **Website:** cfg.eu

 **Social Media:** Follow us on [LinkedIn](#) / [X](#) / [BlueSky](#) for the latest updates.

 **Office Address:** Avenue des Arts 44, 1040 Brussels, Belgium

For media inquiries, speaking engagements, or briefing requests, please contact our communications team at communications@cfg.eu.

Want to stay updated? Subscribe to [our newsletter](#).

Reference List

AI Snake Oil, 'GPT-4 and Professional Benchmarks: The Wrong Answer to the Wrong Question', AI Snake Oil, 2023, available at <https://www.aisnakeoil.com/p/gpt-4-and-professional-benchmarks>.

ARC Prize, 'ARC Prize – What Is ARC-AGI?', ARC Prize, available at <https://arcprize.org/arc>.

Ars Technica, 'OpenAI Announces o3 and o3-mini, Its Next Simulated Reasoning Models', Ars Technica, 2024, available at <https://arstechnica.com/information-technology/2024/12/openai-announces-o3-and-o3-mini-its-next-simulated-reasoning-models/>.

Ansari, S., 'AI Valuation Multiples 2024', Aventis Advisors, 2025, available at <https://aventis-advisors.com/ai-valuation-multiples/>

Bainbridge, W.S. and Roco, M.C., 'Science and technology convergence: with emphasis for nanotechnology-inspired convergence', Journal of Nanoparticle Research, vol. 18, no. 7, 2016, article no. 211, available at <https://link.springer.com/article/10.1007/s11051-016-3520-0>.

Ball, D.W., 'The Era of Predictive AI Is Almost Over', The New Atlantis, no. 77, 2024, pp. 26–35, available at <https://www.thenewatlantis.com/publications/the-era-of-predictive-ai-is-almost-over>.

Baraniuk, C., 'Electricity grids creak as AI demands soar', BBC News, 2024, available at <https://www.bbc.com/news/articles/cj5ll89dy2mo>.

Benaiche, E. (2023) 'AI Washing', Goaland, available at <https://www.goaland.com/blog/AI/ai-washing>.

Bender, E.M. et al., 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?', in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 2021, pp. 610–623, available at https://dl.acm.org/doi/pdf/10.1145/3442188.3445922?utm_source=miragenews&utm_medium=miragenews&utm_campaign=news.

Bick, A., Blandin, A., and Deming, D.J., 'The Rapid Adoption of Generative AI', 2024, available at https://889099f7-c025-4d8a-9e78-9d2a22e8040f.usfiles.com/ugd/889099_fa955a5012dd45abb84dd647ce02de95.pdf.

Choi, J.H., Monahan, A. and Schwarcz, D., 'Lawyering in the Age of Artificial Intelligence', SSRN, 2023, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4626276.

Dafoe, A., 'AI Governance: A Research Agenda', Centre for the Governance of AI, Future of Humanity Institute, University of Oxford, 2018, available at <https://cdn.governance.ai/GovAI-Research-Agenda.pdf>.

'Data Page: The cost of 66 different technologies over time', Our World in Data, 2025, data adapted from Farmer and Lafond, available at <https://ourworldindata.org/grapher/costs-of-66-different-technologies-over-time>.

DeepSeek, 'DeepSeek-R1 Release', DeepSeek API Documentation, 20 January 2025, available at <https://api-docs.deepseek.com/news/news250120>.

Dell'Acqua, F. et al., 'Navigating The Jagged Technological Frontier: Field Experimental Evidence Of The Effects Of AI On Knowledge Worker Productivity And Quality', Harvard Business School, 2023, available at <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>.

Ding, J., Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition, Princeton, NJ, Princeton University Press, 2024, available at <https://press.princeton.edu/books/paperback/9780691260341/technology-and-the-rise-of-great-powers?srsId=AfmBOopELBKgsfqOtTJYOzxeFzxfatALnpKgDfT-I5q4MIIViWaLnGjl>.

Eriksson, M. et al., 'Can We Trust AI Benchmarks? An Interdisciplinary Review of Current Issues in AI Evaluation', ArXiv, 2025, available at <https://arxiv.org/abs/2502.06559>.

Eqvista, 'Price-to-Sales Ratio by Industry', Eqvista, 2025, available at <https://eqvista.com/price-to-sales-ratio-by-industry>.

Evans, O., Chua, J., and Lin, S., 'New, improved multiple-choice TruthfulQA', LessWrong, 2025, available at <https://www.lesswrong.com/posts/Bunfwz6JsNd44kgLT/new-improved-multiple-choice-truthfulqa>.

Fast Company, 'Microsoft's Inflection AI grab cost more than \$1 billion', 2024, available at <https://www.fastcompany.com/91069182/microsoft-inflection-ai-exclusive>.

Financial Times, 'Google DeepMind Duo Share Nobel Chemistry Prize with US Biochemist', Financial Times, 2024, available at <https://www.ft.com/content/4541c07b-f5d8-40bd-b83c-12c0fd662bd9>.

Floridi, L., 'Why the AI Hype is Another Tech Bubble', Philosophy & Technology, Springer, 2024, available at <https://link.springer.com/article/10.1007/s13347-024-00817-w>.

Fodor, J., 'Line Goes Up? Inherent Limitations of Benchmarks for Evaluating Large Language Models', arXiv, 2025, available at <https://arxiv.org/abs/2502.14318>.

Forbes, 'AI Spending to Exceed a Quarter Trillion Next Year', Forbes, 2024, available at <https://www.forbes.com/sites/bethkindig/2024/11/14/ai-spending-to-exceed-a-quarter-trillion-next-year/>.

Fridman, L., 'Dario Amodei: Anthropic CEO On Claude, AGI & The Future Of AI & Humanity | Lex Fridman Podcast #452', YouTube, 2024, <https://www.youtube.com/watch?v=ugvHCXC0mm4>.

Goh, E. et al., 'Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial', JAMA Network Open, 2024, available at <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2825395>.

Goldman Sachs, 'AI Is Showing 'Very Positive' Signs of Eventually Boosting GDP and Productivity', Goldman Sachs, 2024, available at <https://www.goldmansachs.com/insights/articles/AI-is-showing-very-positive-signs-of-boosting-gdp>.

Google, 'Alphabet Q3 Earnings Call: CEO Sundar Pichai's Remarks', Google (blog), 2024, <https://blog.google/inside-google/message-ceo/alphabet-earnings-q3-2024/#full-stack-approach>.

Google, 'Introducing Gemini 2.0 | Our Most Capable AI Model Yet', YouTube, 2024, available at <https://www.youtube.com/watch?v=Fs0t6SdODd8&t=45s>.

Grace, K. et al., 'Thousands of AI Authors on the Future of AI', arXiv, 2024, <https://arxiv.org/abs/2401.02843>.

Hayes, A., 'What Is Enterprise Value-to-Sales (EV/Sales)? How to Calculate', Investopedia, 2021, available at <https://www.investopedia.com/terms/e/enterprisevaluesales.asp>

Hernström, V. et al., 'Screening Performance and Characteristics of Breast Cancer Detected in the Mammography Screening with Artificial Intelligence Trial (MASAI): A Randomised, Controlled, Parallel-Group, Non-Inferiority, Single-Blinded, Screening Accuracy Study', The Lancet, 2025, available at

<https://www.thelancet.com/journals/landig/article/PIIS2589-7500%2824%2900267-X/fulltext>.

Hestness, J. et al. (2017) 'Deep Learning Scaling Is Predictable, Empirically', ArXiv, available at <https://arxiv.org/abs/1712.00409>.

Hinton, G. [@geoffreyhinton], Twitter post, 2023, available at <https://x.com/geoffreyhinton/status/1653687894534504451>.

Hu, K., 'ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note', Reuters, 2023, available at <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>.

Hu, K. and Cai, K., 'OpenAI offers one investor a sweetener that no others are getting', Reuters, 28 September 2024, available at <https://www.reuters.com/technology/artificial-intelligence/openai-sees-116-billion-revenue-next-year-offers-thrive-chance-invest-again-2025-2024-09-28/>.

Hu, K. and Tong, A., 'OpenAI and others seek new path to smarter AI as current methods hit limitations', Reuters, 2024, available at <https://www.reuters.com/technology/artificial-intelligence/openai-rivals-seek-new-path-smarter-ai-current-methods-hit-limitations-2024-11-11/>.

Janků, D., Reddel, F., Graabak, J. and Reddel, M., 'Beyond the AI Hype - A Critical Assessment of AI's Transformative Potential', Centre for Future Generations, 2025, available at: <https://cfg.eu/beyond-the-ai-hype>.

Jayatunga, M.K.P., Xie, W., Ruder, L., Schulze, U., and Meier, C., 'AI in small-molecule drug discovery: a coming wave?', Nature Reviews Drug Discovery, vol. 21, 2022, pp. 175-176, available at <https://www.nature.com/articles/d41573-022-00025-1>.

Jones, N., "'In awe": scientists impressed by latest ChatGPT model o1', Nature, 2024, available at <https://www.nature.com/articles/d41586-024-03169-9>.

Juijn, D., 'No, Deep Learning Is (probably) Not Hitting A Wall', Centre for Future Generations, 2024, available at <https://cfg.eu/no-deep-learning-is-probably-not-hitting-a-wall/>.

Krishnan, N., 'OpenAI's O3: A New Frontier in AI Reasoning Models', Towards AI, 2024, <https://pub.towardsai.net/openais-o3-a-new-frontier-in-ai-reasoning-models-a02999246ffe>.

LeCun, Y., 'About the raging debate regarding the future of AI', Facebook post, 2024, available at <https://www.facebook.com/yann.lecun/posts/10158256523332143>.

LeGris, S., Vong, W. K., Lake, B. M., and Gureckis, T. M., 'H-ARC: A Robust Estimate of Human Performance on the Abstraction and Reasoning Corpus Benchmark', arXiv, 2024, <https://arxiv.org/pdf/2409.01374>.

Lin, S., Hilton, J., and Evans, O., 'TruthfulQA: Measuring How Models Mimic Human Falsehoods', arXiv, 2021, <https://arxiv.org/abs/2109.07958>.

Lin, T., 'Measuring Automated Kernel Engineering', METR, 2025, <https://metr.org/blog/2025-02-14-measuring-automated-kernel-engineering/>.

Marcus, G., 'Folks, Game Over....', X, 2024, available at <https://x.com/GaryMarcus/status/1855382564015689959>.

Marcus, G., 'Deep Learning Is Hitting a Wall', Nautilus, 2024, available at <https://nautil.us/deep-learning-is-hitting-a-wall-238440/>.

Marcus, G., 'What If Generative AI Turned Out to be a Dud?', Communications of the ACM, vol. 66, no. 3, 2023, available at <https://cacm.acm.org/blogcacm/what-if-generative-ai-turned-out-to-be-a-dud/>.

McClain, C., Kennedy, B., Gottfried, J., Anderson, M., and Pasquini, G., 'How the U.S. Public and AI Experts View Artificial Intelligence', Pew Research Center, 2025, available at <https://www.pewresearch.org/2025/04/03/artificial-intelligence-in-daily-life-views-and-experiences/>.

McKinsey, 'The State of AI in 2024', McKinsey & Company, 2024, available at <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

Metaculus, 'When will AIs program programs that can program AIs?', Metaculus, n.d., <https://www.metaculus.com/questions/406/date-ais-capable-of-developing-ai-software/>.

Metz, R. et al., 'OpenAI, Google and Anthropic Are Struggling to Build More Advanced AI', Bloomberg, 2024, available at <https://www.bloomberg.com/news/articles/2024-11-13/openai-google-and-anthropic-are-struggling-to-build-more-advanced-ai>.

Moravec, H., *Mind Children*, Harvard University Press, 1988.

National History Education Clearinghouse, 'Innovation and Technology in the 19th Century', Teachinghistory.org, n.d., available at <https://teachinghistory.org/history-content/ask-a-historian/24470>.

National Research Council, Global Affairs, and Panel on the Impact of Information Technology on the Future of the Research University, Preparing for the Revolution: Information Technology and the Future of the Research University, Washington, DC, National Academies Press, 2002, available at <https://nap.nationalacademies.org/read/10545/chapter/3#7>.

NYU Stern, 'Price to Sales Ratios', 2025, available at https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/psdata.html.

Olson, P., 'Nearly Half Of All 'AI Startups' Are Cashing In On Hype', Forbes, 4 March 2019, available at <https://www.forbes.com/sites/parmyolson/2019/03/04/nearly-half-of-all-ai-startups-are-cashing-in-on-hype/>.

OpenAI Developer Community, 'Incorrect Count of 'r' Characters in the Word 'strawberry' – Use Cases and Examples', OpenAI Developer Community, 2024, available at <https://community.openai.com/t/incorrect-count-of-r-characters-in-the-word-strawberry/829618>.

'OpenAI o3 Breakthrough High Score On ARC-AGI-pub', ARC Prize, 2024, <https://arcprize.org/blog/oai-o3-pub-breakthrough>.

OpenAI, 'Pricing', OpenAI, 2023, available at <https://web.archive.org/web/20230401144811/https://openai.com/pricing>.

Our World in Data, 'Cost of Sequencing a Full Human Genome', Our World in Data, 2023, available at <https://ourworldindata.org/grapher/cost-of-sequencing-a-full-human-genome>.

Our World in Data, 'Share of United States Households Using Specific Technologies', Our World in Data, n.d., available at <https://ourworldindata.org/grapher/technology-adoption-by-households-in-the-united-states>.

Our World in Data, 'Solar PV Prices', Our World in Data, 2024, available at <https://ourworldindata.org/grapher/solar-pv-prices>.

Pacchiardi, L., Tesic, M., Cheke, L.G., and Hernández-Orallo, J., 'Leaving the barn door open for Clever Hans: Simple features predict LLM benchmark answers', arXiv, 2024, available at <https://arxiv.org/abs/2410.11672>.

Parasnis, S., 'Bank of China announces 1 Trillion Yuan AI Industry investment', Medianama, 2025, available at <https://www.medianama.com/2025/01/223-bank-of-china-announces-1-trillion-yuan-ai-industry-investment/>.

Pavlus, J., 'Machines Beat Humans on a Reading Test. But Do They Understand?', Quanta Magazine, 17 October 2019, available at <https://www.quantamagazine.org/machines-beat-humans-on-a-reading-test-but-do-they-understand-20191017/>.

Pew Research Center, 'Americans Increasingly Using ChatGPT, but Few Trust Its 2024 Election Information', Pew Research Center, 2024, available at <https://www.pewresearch.org/short-reads/2024/03/26/americans-use-of-chatgpt-is-ticking-up-but-few-trust-its-election-information/>.

Pillay, T., 'How OpenAI's Sam Altman Is Thinking About AGI and Superintelligence in 2025', TIME, 2025, <https://time.com/7205596/sam-altman-superintelligence-agi/>.

Piquard, A., 'Chatbots Are Like Parrots, They Repeat Without Understanding', Le Monde, 2024, available at https://www.lemonde.fr/en/economy/article/2024/10/07/chatbots-are-like-parrots-they-repeat-without-understanding_6728523_19.html.

Priestley, M., Sluckin, T.J. and Tiropanis, T., 'Innovation on the Web: the end of the S-curve?', Internet Histories, vol. 4, 2020, available at https://www.researchgate.net/publication/340525808_Innovation_on_the_Web_the_end_of_the_S-curve.

Rao, A. et al., 'Assessing the Utility of ChatGPT Throughout the Entire Clinical Workflow: Development and Usability Study', Journal of Medical Internet Research, 2023, available at <https://www.jmir.org/2023/1/e48659/>.

Reddit, 'ChatGPT getting basic math wrong', 2023, available at https://www.reddit.com/r/ChatGPT/comments/103bw12/chat_gpt_getting_basic_maths_wrong/.

Reuters, '400 Million Weekly Users', Reuters, 2025, available at <https://www.reuters.com/technology/artificial-intelligence/openais-weekly-active-users-surpass-400-million-2025-02-20/>.

Seeking Alpha, 'Microsoft Corporation (MSFT) Q1 2025 Earnings Call Transcript', Seeking Alpha, 2025, available at <https://seekingalpha.com/article/4731223-microsoft-corporation-msft-q1-2025-earnings-call-transcript>.

Sevilla, J. et al., 'Can AI Scaling Continue Through 2030?', Epoch AI, 2024, available at <https://epoch.ai/blog/can-ai-scaling-continue-through-2030>.

Silva, J., Sherman, N., and Rahman-Jones, I., 'Stargate: Tech giants announce AI plan worth up to \$500bn', BBC News, 2024, available at <https://www.bbc.com/news/articles/cy4m84d2xz2o>.

Stanford HAI, 'HAI AI-Index-Report-2024', Stanford HAI, 2024, available at https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf.

Toner-Rodgers, A., 'Artificial Intelligence, Scientific Discovery, and Product Innovation', NBER Conference Papers, 2024, available at https://conference.nber.org/conf_papers/f210475.pdf.

Tsarathustra, 'Dane Vahey of OpenAI Says the Cost...', X, 2024, available at <https://x.com/tsarnick/status/1838401544557072751>.

Tuladhar, S.K., 'Two things ChatGPT can get wrong: math and facts. Why?', Medium, 2023, available at <https://medium.com/@k.sunman91/two-things-chatgpt-can-get-wrong-math-and-facts-why-b40e50dc68cd>.

UK Government, 'International AI Safety Report', UK Government, 2025, https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.

U.S. Food and Drug Administration (FDA), 'Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices', FDA, 2024, available at <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-ai-ml-enabled-medical-devices>.

Varanasi, L., 'Here's How Far We Are From AGI, According to the People Developing It', Business Insider, 2024, available at <https://www.businessinsider.com/agi-predictions-sam-altman-dario-amodei-geoffrey-hinton-demis-hassabis-2024-11>.

Vectara, 'Hallucination Leaderboard', GitHub, n.d., <https://github.com/vectara/hallucination-leaderboard>.

The Wall Street Journal, 'Struggling AI Startups Look for a Bailout From Big Tech', 2024, available at <https://www.wsj.com/tech/ai/struggling-ai-startups-look-for-a-bailout-from-big-tech-3e635927>.

Wei, J. et al., 'Emergent Abilities of Large Language Models', arXiv, 2022, available at <https://arxiv.org/abs/2206.07682>.

Wijk, H. et al., 'Re-bench: Evaluating Frontier AI R&D Capabilities Of Language Model Agents Against Human Experts', METR, 2024, available at <https://metr.org/blog/2024-11-22-evaluating-r-d-capabilities-of-llms/>.

Wikipedia, 'AI washing', 2025, available at https://en.wikipedia.org/wiki/AI_washing.

Wold, J.W., 'Von der Leyen Launches World's Largest Public-Private Partnership To Win AI Race', EURACTIV, 2025, available at <https://www.euractiv.com/section/tech/news/von-der-leyen-launches-worlds-largest-public-private-partnership-to-win-ai-race/>.

Yee, L., Chui, M., and Roberts, R., with Xu, S., 'Why agents are the next frontier of generative AI', McKinsey Digital, 2024, <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai>.

YouGov, 'Do Americans Think AI Will Have a Positive or Negative Impact on Society?', YouGov, 2025, available at <https://today.yougov.com/technology/articles/51368-do-americans-think-ai-will-have-positive-or-negative-impact-society-artificial-intelligence-poll>.

