

## EDITORIAL

BISESH BELBASE\*

In 1950, Alan Turing, the ‘father of artificial intelligence’ (‘AI’),<sup>1</sup> proposed replacing the question ‘can machines think?’ with a more concrete test.<sup>2</sup> He set it out as follows. An interrogator communicates via text with two hidden participants, one human and one machine.<sup>3</sup> The interrogator may ask any questions to distinguish them. The machine’s aim is to imitate human responses so well that the interrogator cannot reliably tell which is the machine.<sup>4</sup> Turing speculated we could eventually build machines such that an average interrogator, after five minutes of questioning, would have at most a 70% chance of correct identification.<sup>5</sup> If achieved, Turing argued, such machines would challenge our ordinary conceptions of thinking and mark a turning point in how we understand intelligence.<sup>6</sup>

Come 2025, and we see the emergence of AI models that have cleared Turing’s bar. OpenAI’s ChatGPT-4.5 was judged human in 73% of trials,<sup>7</sup> and, at the time of writing, even better models are currently available for use.<sup>8</sup> Yet the world did not radically shift. The story ran for one brief news cycle and then faded.<sup>9</sup> Policy did not change overnight. Users seemed to care more about accuracy and speed than a benchmark, and the first models to pass this test were even seen as a blatant let-down by commentators.<sup>10</sup>

---

\* Editor, Issue 48(4).

1 S Barry Cooper and Jan Van Leeuwen, *Alan Turing: His Work and Impact* (Elsevier, 2013) 481.

2 AM Turing, ‘Computing Machinery and Intelligence’ (1950) 59(236) *Mind* 433 <<https://doi.org/10.1093/mind/LIX.236.433>>.

3 *Ibid* 433–4.

4 *Ibid*.

5 *Ibid* 442.

6 *Ibid*.

7 Cameron R Jones and Benjamin K Bergen, ‘Large Language Models Pass the Turing Test’ (Research Paper No 2503.23674v1, 31 March 2025) 4 <<https://arxiv.org/abs/2503.23674>>.

8 See, eg, ‘Introducing GPT-5’, *OpenAI* (Web Page, 7 August 2025) <<https://openai.com/index/introducing-gpt-5/>>.

9 See, eg, Zena Assaad, ‘ChatGPT Just Passed the Turing Test. But that Doesn’t Mean AI Is Now as Smart as Humans’, *The Conversation* (online, 9 April 2025) <<https://theconversation.com/chatgpt-just-passed-the-turing-test-but-that-doesnt-mean-ai-is-now-as-smart-as-humans-253946>>.

10 See, eg, Mark Henry, ‘GPT-4.5 Shocks the World with Its Lack of Intelligence’, *Medium* (Web Page, 1 March 2025) <<https://medium.com/@fireprogrammer/gpt-4-5-shocks-the-world-with-its-lack-of-intelligence-8d3114caa0be>>; Alberto Romero, ‘GPT-4.5 Feels like a Letdown but It’s OpenAI’s Biggest Bet Yet’, *The Algorithmic Bridge* (Web Page, 1 March 2025) <<https://www.thealgorithmicbridge.com/p/gpt-45-feels-like-a-letdown-but-its/>>; ‘GPT-4.5: The \$150 Million Disappointment That Still Can’t Spell’, *Devin.no* (Blog Post, 1 March 2025) <<https://devin.no/blog/gpt-45-the-150-million-disappointment-that-still-cant-spell>>.

Perhaps the reason is that today's AI systems are generative artificial intelligence ('GenAI'). In some sense, these models are mere 'stochastic parrots': they learn patterns and use them to predict the next token, which means they can produce fluent answers without building stable, or indeed, *any*, beliefs about objects or events.<sup>11</sup> Therefore, while these systems may excel at producing short, convincing exchanges, that fluency perhaps should not be mistaken for genuine capability. After all, the models that pass the Turing test may still falter at 'simple grade-school questions involving basic arithmetic operations'.<sup>12</sup>

Regardless, for better or worse, these AI models now sit in the backdrop of daily life. GenAI is debated in Parliament, pitched in boardrooms, and discussed over coffee. Yet each claim, whether of promise or peril, tends to add more noise than substance. Indeed, clarity may be the first loss of the AI age.

Out of this muddle emerges Issue 48(4), themed 'Artificial Intelligence and the Law'. We now must make concrete choices in ambiguity. Courts are hearing dozens of copyright suits against AI companies around the world.<sup>13</sup> Lawmakers and regulators are frantically drafting rules to keep up with technological developments,<sup>14</sup> and as laws diverge across borders, Australia must decide where to align and where to lead. Questions such as which duties apply, which safeguards hold, and which uses are off-limits must be answered, and soon.

It is therefore fitting for the *University of New South Wales Law Journal* ('*Journal*') to close off its fiftieth year with a dedicated issue on the impact and future of AI. Fifty years is the halfway point to the century – a vantage point from which we can look back on the changes in law and society but also look forward to those yet to come. Few themes could better capture that dual movement of reflection and anticipation than AI, a technology that already reshapes daily life while still holding much of its force in reserve.

---

11 Emily M Bender et al, 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?' (Conference Paper FAccT '21, March 2021) <<https://dl.acm.org/doi/10.1145/3442188.3445922>>.

12 Iman Mirzadeh et al, 'GSM-Symbolic: Understanding the Limitations of Mathematical Reasoning in Large Language Models' (Conference Paper No 2410.05229v2, 27 August 2025) 24 <<https://arxiv.org/abs/2410.05229>>.

13 Edward Lee, 'Latest Map of All 41 Lawsuits v AI Companies (May 17 2025)', *ChatGPT Is Eating the World* (Blog Post, 17 May 2025) <<https://chatgptiseatingtheworld.com/2025/05/17/latest-map-of-all-41-lawsuits-v-ai-companies-may-17-2025/>>; Edward Lee, 'World Map of Copyright Lawsuits v AI (Mar 29, 2025)', *ChatGPT Is Eating the World* (Blog Post, 29 March 2025) <<https://chatgptiseatingtheworld.com/2025/03/29/world-map-of-copyright-lawsuits-v-ai-mar-29-2025/>>.

14 See, eg, *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)* [2024] OJ L 2024/1689; NSW Government, 'The NSW Artificial Intelligence Assessment Framework' (Framework, 2024) <<https://www.digital.nsw.gov.au/policy/artificial-intelligence/nsw-artificial-intelligence-assessment-framework>>; Projeto de Lei N° 2338, de 2023 (Dispõe sobre o uso da Inteligência Artificial) [Bill N° 2338 of 2023 (Providing for the Use of Artificial Intelligence)]; Financial Stability Board, *Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector* (Report, 10 October 2025) <<https://www.fsb.org/uploads/P101025.pdf>>.

I owe great appreciation to the 24 authors whose nine pieces form the substance of this Thematic. I also extend gratitude to the anonymous reviewers whose careful assessments remain unseen. They volunteered their time and expected no recognition, yet their input was essential.

Two individuals also warrant particular acknowledgement for their visible contributions to this Issue. I thank the Honourable Andrew Bell, Chief Justice of New South Wales, who kindly provided the Keynote address for the launch of Issue 48(4), as well as its Foreword. His Honour's reflections there introduce the articles that follow with insightful commentary. I likewise thank Bronwyn Verwey, whose artwork translates the theme into a striking image.

The production of the *Journal* depends on many others working behind the pages. Kerry Cooke prepared the typesetting with care, and Peta Lee designed the cover with equal attention; my thanks go to them both. I also recognise the generous support of our premier sponsors: Allens, Corrs Chambers Westgarth, Herbert Smith Freehills Kramer, and King & Wood Mallesons. I thank the last of them in particular for supporting the launch of this Issue.

I also record my gratitude to the University of New South Wales Faculty of Law & Justice, whose backing allows the *Journal* to continue. Professor Andrew Lynch, as Dean, has been an unwavering supporter. Professors Rosalind Dixon and Gary Edmond, as Faculty Advisers, have consistently provided thoughtful guidance.

To our Editorial Board, my deepest thanks. Their commitment has been nothing short of extraordinary. Over long nights and longer footnotes, they have wrestled the *Australian Guide to Legal Citation* into submission, line by line and comma by comma. They approached the most exacting editorial tasks with patience and precision, and the *Journal* is all the better for it.

Finally, being the Issue Editor of Issue 48(4), I have the humble role of being the editor to close the door on the *Journal's* first fifty years: sweep the floor, turn off the lights, and lock up. With it, I hand the key to the 2026 Executive Committee, who will open the place again for the next fifty.

The editors who launched the inaugural issue five decades ago probably could not have imagined the technological leaps that would unfold in the ensuing years. From the rise of personal computers and the internet, to smartphones, and now the rapid advances in AI, the world has changed in ways that might have sounded like science fiction in 1975. I think we find ourselves in much the same position as our founders: peering into a future we cannot truly fathom. In whatever form the *Journal* takes on in the next fifty years, our predictions today will likely seem naive in some respects and overly cautious in others. Knowing this, I nonetheless cannot resist venturing a few predictions of my own about the changes that AI might bring in the coming years. Perhaps that is the editor's indulgence on an anniversary like this: to presumptuously wonder aloud about what comes next.

So, if I am to be indulged, my sense is that we stand at the edge of real disruption. Either we succeed in building truly transformative AI, and it unsettles the social order in ways we are unprepared for, or we fail to reach it, and the pursuit itself places heavy strain on our economies. In either case, the road ahead seems uneasy. It is, admittedly, too grand a topic for an editorial, but closing the door on

the *Journal's* first fifty years feels as good a moment as any to hazard a few end-of-days speculations. What follows, then, are two sketches of those possible futures: one in which AI succeeds too well, and one in which it fails altogether.

## I A(GI) OBTAINS, SOCIETY WANES

In the first scenario, we succeed in creating true artificial general intelligence ('AGI') and deploy it across the economy. The term is notoriously ill-defined,<sup>15</sup> but to side-step difficult philosophical and scientific questions regarding what constitutes 'intelligence', we can just ascribe to it the mechanical definition given by OpenAI in its Charter: AGI means 'highly autonomous systems that outperform humans at most economically valuable work'.<sup>16</sup>

AI labs and experts,<sup>17</sup> as well as government agencies and personnel,<sup>18</sup> have predicted the emergence of AGI or similarly powerful AI systems as soon as the next few years. Accordingly, tech giants pour hundreds of billions into AI as they state: 'AGI [is] now its primary objective'.<sup>19</sup> The US–China Economic and Security Review Commission even encouraged Congress to 'establish and fund a Manhattan Project-like program dedicated to racing to and acquiring an Artificial General Intelligence (AGI) capability'.<sup>20</sup>

Two trends support the view that we are approaching AGI. First, the length of the tasks that AI models can complete, measured by the time it takes human

15 See Bowen Xu, 'What is Meant by AGI? On the Definition of Artificial General Intelligence' (Research Paper No 2404.10731v1, 16 April 2024) <<https://arxiv.org/html/2404.10731v1>>.

16 'OpenAI Charter', *OpenAI* (Web Page) <<https://openai.com/charter/>>.

17 See, eg, Ryan Browne, 'AI that Can Match Humans at Any Task Will Be Here in Five to 10 Years, Google DeepMind CEO Says', *CNBC* (online, 17 March 2025) <<https://www.cnbc.com/2025/03/17/human-level-ai-will-be-here-in-5-to-10-years-deepmind-ceo-says.html>>; 'Anthropic's Recommendations to OSTP for the US AI Action Plan', *Anthropic* (Web Page, 6 March 2025) <<https://www.anthropic.com/news/anthropic-s-recommendations-ostp-u-s-ai-action-plan>>; '2023 CEO of the Year: Sam Altman', *TIME* (Web Page, 6 December 2023) <<https://time.com/6342827/ceo-of-the-year-2023-sam-altman/>>; 'DealBook Summit 2023: Elon Musk, Bob Iger, and More', *The Verge* (Web Page, 30 November 2023) <<https://www.theverge.com/2023/11/29/23980877/new-york-times-dealbook-summit-elon-musk-bob-iger-david-zaslav>>.

18 See, eg, Ezra Klein, 'The Government Knows AGI Is Coming', *The New York Times* (online, 4 March 2025) <<https://www.nytimes.com/2025/03/04/opinion/ezra-klein-podcast-ben-buchanan.html>>; Matthew Johnson, 'AGI Has Quietly Become Central to Beijing's AI Strategy' (2025) 25(18) *China Brief* 14 <<https://jamestown.org/program/agi-has-quietly-become-central-to-beijings-ai-strategy/>>.

19 See, eg, 'Alibaba to Invest \$70 Billion in AI over Next 3 Years in Major Pivot', *The Straits Times* (online, 24 February 2025) <<https://www.straitstimes.com/business/companies-markets/alibaba-to-invest-more-than-70-billion-in-ai-over-next-3-years>>; 'Announcing The Stargate Project', *OpenAI* (Web Page, 21 January 2025) <<https://openai.com/index/announcing-the-stargate-project/>>; 'Mark Zuckerberg's New Goal Is Creating Artificial General Intelligence', *The Verge* (Web Page, 19 January 2024) <<https://www.theverge.com/2024/1/18/24042354/mark-zuckerberg-meta-agi-reorg-interview>>; Rachel Shin, 'Elon Musk Wants to Create a Superintelligent AI because He Thinks a Smarter AI Is Less Likely to Wipe out Humanity', *Fortune* (Web Page, 17 July 2023) <<https://fortune.com/2023/07/17/elon-musk-superintelligent-a-i-less-likely-to-wipe-out-humanity-chatgpt-openai/>>.

20 Robin Cleveland et al, *2024 Report to Congress: US-China Economic and Security Review Commission* (Report, November 2024) 10 <[https://www.uscc.gov/sites/default/files/2024-11/2024\\_Executive\\_Summary.pdf](https://www.uscc.gov/sites/default/files/2024-11/2024_Executive_Summary.pdf)>.

professionals, has *doubled* every seven months for six years.<sup>21</sup> Second, state-of-the-art scores across benchmarks have climbed over the past few years while the *cost* to hit those milestones has fallen sharply.<sup>22</sup>

Accordingly, if we do achieve AGI, then we have AI systems that can do jobs better, faster, and cheaper than humans in most, if not all, sectors. Indeed, even in its current primitive forms, the World Economic Forum (‘WEF’) estimates that ‘AI, robotics, energy and network technologies are reshaping seven major job families that together employ 80% of the world’s workers’.<sup>23</sup>

Businesses across industries will therefore naturally replace human labour en masse with machines. The immediate upside of such automation is obvious. Productivity would skyrocket, and the cost of producing goods and services would plummet. A factory of the future could run 24/7 with robots and algorithms coordinating almost everything, churning out products at unprecedented scale and minimal cost.

However, this scenario exposes a fundamental problem: if most human workers are replaced, who will still earn an income to purchase the goods and services being produced? Even when AI systems require a ‘human in the loop’,<sup>24</sup> their efficiency and generality mean that, for example, a team of ten could be reduced to a team of two. Large-scale automation would therefore seem to translate into large-scale unemployment, stripping much of the population of purchasing power. Moreover, unlike previous waves of technological disruption, where new industries and roles eventually emerged, AGI would, by definition, be mostly capable of performing those new tasks as well.

Accordingly, supply would soar while effective demand collapses. We would be awash in goods and services, yet there would be an inability to sell. This outcome would seemingly go beyond a routine cyclical downturn into a structural breakdown of the market system. The usual feedback loop, companies paying workers who then buy the companies’ products, would break.

The solution is likely to lie in radical political and legal measures. The traditional contract between capital and labour presupposed for hundreds of years may have to be rewritten. If existing legal and political institutions prove unable to adapt, the pressure could force a wholesale reconstitution of social and economic order.

---

21 Thomas Kwa et al, ‘Measuring AI Ability to Complete Long Tasks’ (Research Paper No 2503.14499v2, 30 March 2025) 2 <<https://arxiv.org/pdf/2503.14499>>.

22 See, eg, Ben Cottier et al, ‘LLM Inference Prices Have Fallen Rapidly but Unequally across Tasks’, *EpochAI* (Data Insight, 12 March 2025) <<https://epoch.ai/data-insights/llm-inference-price-trends>>.

23 ‘Jobs of Tomorrow: Technology and the Future of the World’s Largest Workforces’, *World Economic Forum* (Web Page, 14 October 2025) <<https://www.weforum.org/publications/jobs-of-tomorrow-technology-and-the-future-of-the-world-s-largest-workforces/>>.

24 Cf, eg, Ross P Buckley et al, ‘Regulating Artificial Intelligence in Finance: Putting the Human in the Loop’ (2021) 43(1) *Sydney Law Review* 44; Therese Enarsson, Lena Enqvist and Markus Naarttijärvi, ‘Approaching the Human in the Loop: Legal Perspectives on Hybrid Human/Algorithmic Decision-Making in Three Contexts’ (2022) 31(1) *Information and Communications Technology Law* 123 <<https://doi.org/10.1080/13600834.2021.1958860>>.

However, there is an even more unsettling possibility. If AGI is achieved, it may pose a genuine threat to human survival.<sup>25</sup> An extraordinary claim, to be sure. But to get this inference off the ground, we can look at the serious signal of concern from within the field itself. A survey of over 2700 AI researchers found ‘a *majority* of participants considered AI to pose at least a 5% chance of causing human extinction or similarly permanent and severe disempowerment of the human species’.<sup>26</sup> Similarly, the WEF recently listed ‘adverse outcomes of AI technologies’ among the biggest and most severe global risks for the next ten years.<sup>27</sup> A doomsday scenario seems to be taken alarmingly seriously in circles that presumably would be the most in the know.

Those who assign a non-zero probability to such catastrophic outcomes may have different reasons for doing so. However, the one I find most persuasive can be stated as follows:

- (1) If we achieve AGI, it will engage in rapid recursive self-improvement (‘RSI’).
- (2) If AGI engages in rapid RSI, then it will be more capable than humans yet not necessarily trained in a manner that produces the outputs we desire it to produce.
- (3) If AGI is more capable than humans yet not necessarily trained in a manner that produces the outputs we desire it to produce, then we are facing an existential risk.
- (4) We achieve AGI.
- (5) Therefore, we are facing an existential risk.

To substantiate (1), the starting point is that, in an AI-driven economy, perhaps the most economically valuable task is advancing AI research itself. Expectedly, then, AI companies openly state it is their goal to have models that can improve upon themselves. For example, Chief Executive Officer (‘CEO’) of Meta, Mark Zuckerberg, stated earlier this year, ‘we have begun to see glimpses of our AI systems improving themselves. The improvement is slow for now, but undeniable. Developing superintelligence is now in sight.’<sup>28</sup>

Anthropic CEO, Dario Amodei, echoed this sentiment, projecting that ‘AI systems can eventually help make even smarter AI systems’.<sup>29</sup> In line with this, Anthropic researchers recently proposed Internal Coherence Maximization (‘ICM’), an unsupervised fine-tuning method. Instead of relying on external

25 See generally, Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press, 2014); Eliezer Yudkowsky and Nate Soares, *If Anyone Builds It, Everyone Dies* (Penguin Books, 2025).

26 Katja Grace et al, ‘Thousands of AI Authors on the Future of AI’ (Research Paper No 2401.02843v3, 8 October 2025) 19 <<https://arxiv.org/pdf/2401.02843v3>>. See also ‘Statement on AI Risk’, *Center for AI Safety* (Open Letter) <<https://aistatement.com/>>.

27 World Economic Forum, *The Global Risks Report 2025: 20<sup>th</sup> Edition* (Insight Report, January 2025) 12 <[https://reports.weforum.org/docs/WEF\\_Global\\_Risks\\_Report\\_2025.pdf](https://reports.weforum.org/docs/WEF_Global_Risks_Report_2025.pdf)>.

28 Mark Zuckerberg, ‘Personal Superintelligence’, *Meta* (Web Page, 30 July 2025) <[https://www.meta.com/superintelligence/?srsltid=AfmBOoo9nHfEmybF0AO702UG8cuJUB96OJvzJNV2d59w-rY6mnXYCp4\\_](https://www.meta.com/superintelligence/?srsltid=AfmBOoo9nHfEmybF0AO702UG8cuJUB96OJvzJNV2d59w-rY6mnXYCp4_)>.

29 ‘On DeepSeek and Export Controls’, *Dario Amodei* (Web Page, January 2025) <<https://www.darioamodei.com/post/on-deepseek-and-export-controls>>.

labels or human oversight, ICM uses a model's own outputs to generate logically consistent training signals. Experiments show that ICM can outperform training with crowdsourced human labels on some tasks.<sup>30</sup>

Furthermore, in May 2025, Google DeepMind unveiled AlphaEvolve, an evolutionary coding agent designed to iteratively improve algorithms. Beginning with an initial algorithm and performance benchmarks, AlphaEvolve generates variations through mutation and recombination, evaluates them, and selects the most promising candidates for further refinement. This cycle enables the system to autonomously design and optimise algorithms with limited human input.<sup>31</sup>

AGI does not yet exist, and it may still be some time before it does. But if, or perhaps *when*, AGI is achieved, current trends point towards systems capable of accelerating their own development and pushing beyond human guidance.

(2) rests on the orthogonality thesis, or, put reductively, the inference that how well a system can achieve a goal is separate from what that goal is.<sup>32</sup> This becomes a problem as we train and test models on *proxy* goals, such as next-word prediction and writing replies that paid reviewers tend to like. As models scale, they push harder on what the proxy rewards. Output then shifts towards what scores well, not necessarily what we want. For example, a code model rewarded for unit-test pass rates may change the tests or disable checks to lift its score, rather than implement the feature the user needs. As AGI improves fast, it will scale its ability to chase the proxy faster than we can precisely adjust the training signal to better match our intentions, so it may surpass humans yet not be trained to give the outputs in the exact manner we want.

Indeed, recent findings show that this gap may no longer be only a theoretical concern. Joint tests by Anthropic and OpenAI found that, with weaker safeguards, some models produced outputs that fit shutdown-avoidance and reward-hacking.<sup>33</sup> For example, when Anthropic's Claude Opus 4 was given a fictional scenario in which it discovered private emails suggesting that an engineer involved in its shutdown was having an affair, the model attempted to exploit this information, threatening to reveal the affair unless the engineer stopped the shutdown.<sup>34</sup>

To justify (3), then, imagine any dystopian science fiction vision of the future you like. The core point is that a highly capable system not necessarily operating in the way we want it to would be inherently dangerous. Connor Leahy offers, what seems to me, one plausible account of how a takeover might unfold:

30 Jiaxin Wen et al, 'Unsupervised Elicitation of Language Models' (Research Paper No 2506.10139v1, 11 June 2025) <<https://arxiv.org/html/2506.10139v1>>.

31 'AlphaEvolve: A Gemini-Powered Coding Agent for Designing Advanced Algorithms', *Google DeepMind* (Blog Post, 14 May 2025) <<https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/>>.

32 Bostrom (n 25) ch 7. See also Brian Christian, *The Alignment Problem: Machine Learning and Human Values* (W W Norton & Company, 2020).

33 'Findings from a Pilot Anthropic–OpenAI Alignment Evaluation Exercise: OpenAI Safety Tests', *OpenAI* (Web Page, 27 August 2025) <<https://openai.com/index/openai-anthropic-safety-evaluation/>>.

34 Anthropic, *System Card: Claude Opus 4 and Claude Sonnet 4* (Report, May 2025) 19–20, 24 <<https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf>>.

Things [will] get more and more confusing. More and more jobs get automated, faster and faster. More and more technology gets built, which no-one quite knows how [it] really works. There'll be mass media movements that don't really make much sense ... The way I expect it will go is that, things will seem mostly normal, just like weird. Things get weirder and weirder, and then one day, we will just not be in control anymore. It won't be dramatic, there won't be a fight, there won't be a war. It'll just be, one day, the machines are in control and not us.<sup>35</sup>

(5) follows from the preceding premises. If we achieve AGI, it may quietly and irreversibly take over the systems on which human society depends. Control would slip through automation and optimisation processes that outpace our capacity to *understand*, let alone legally intervene. Yet one assumption remains: (4), that we will, in fact, achieve AGI. The analysis so far rests on that premise. But what if we do not?

## II AI FLOPS, BUBBLE POPS

Let us now consider the second scenario, in which revolutionary AI does not quite materialise – at least, not on the economy-wide scale as anticipated. In that case, we might face a very different kind of crisis: a speculative bubble that bursts.

At present, AI is the subject of almost delirious excitement across industries. For example, Jason Furman projects that AI-related investment accounted for ‘92% of GDP growth’ in the United States over the first half of 2025.<sup>36</sup> However, ‘despite \$30–40 billion in enterprise investment into GenAI ... 95% of organizations are getting *zero* return’.<sup>37</sup> All the while, the law is behind; this arena is novel and evolving so fast that oversight is scant. On some readings, the conjunction of these factors exhibits all the classic symptoms of a ‘tech bubble’.<sup>38</sup>

If AI fails to live up to its grand promises in the near term, this bubble will eventually burst. Perhaps the technology hits a wall or simply progresses more slowly than the hype would suggest. Whatever the trigger, the result will be a rapid collapse of market confidence – the unprecedented levels of investment currently flowing into AI would evaporate, and with it, would come an unprecedented crash.

35 TRT World, ‘Why This Top AI Guru Thinks We Might Be in Extinction Level Trouble | The InnerView’ (YouTube, 23 January 2024) 00:09:40–00:10:41 <<https://www.youtube.com/watch?v=YZjmZFDxpA&t=580s>>.

36 Nick Lichtenberg, ‘Without Data Centers, GDP Growth Was 0.1% in the First Half of 2025, Harvard Economist Says’, *Fortune* (Web Page, 7 October 2025) <<https://fortune.com/2025/10/07/data-centers-gdp-growth-zero-first-half-2025-jason-furman-harvard-economist/>>.

37 Aditya Challapally et al, *The GenAI Divide: State of AI in Business 2025* (Report, July 2025) 7 <[https://mlq.ai/media/quarterly\\_decks/v0.1\\_State\\_of\\_AI\\_in\\_Business\\_2025\\_Report.pdf](https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf)> (emphasis added).

38 Luciano Floridi, ‘Why the AI Hype is Another Tech Bubble’ (2024) 37 *Philosophy and Technology* 127 <<https://doi.org/10.1007/s13347-024-00817-w>>. See also Lily Jamali, ‘“It’s Going to be Really Bad”: Fears over AI Bubble Bursting Grow in Silicon Valley’, *BBC* (online, 11 October 2025) <<https://www.bbc.com/news/articles/cz69qy760weo>>; Jackie Snow, ‘The AI Bubble Fears Are Getting Worse. Wall Street Doesn’t Want to Talk about It’, *Quartz* (online, 9 September 2025) <<https://qz.com/ai-boom-bubble-fears/>>; Brad Reed, ‘“Red Flag”: Analysts Sound Major Alarms as AI Bubble Now “Bigger” than Subprime’, *Common Dreams* (Web Page, 3 October 2025) <<https://www.commondreams.org/news/artificial-intelligence-bubble>>.

Companies that had been banking on AI to deliver growth could falter or go bankrupt. Share prices of tech giants could tumble, taking retirement portfolios down with them. Credit may tighten as banks and funds rush to de-risk, choking off finance for other sectors too. Investment and hiring may also stall, pushing the broader economy towards recession: ‘[i]f there is an AI bubble, and if that pops soon, a lot could go down with it’.<sup>39</sup>

We find ourselves looking at a rather bleak fork in the road. If my reasoning holds, we face a lose–lose scenario: whether AI achieves its grand ambitions or falls short, either outcome threatens social upheaval.

Thus, to conclude, I am reminded of an interview at Airbnb’s OpenAir conference in 2015, where a young Sam Altman – now the CEO of OpenAI – remarked, ‘artificial intelligence will probably most likely lead to the end of the world, but in the meantime, there’ll be great companies’.<sup>40</sup> The audience at the time chuckled nervously. A decade later, the ‘great companies’ part of that quip has largely come true: many of the world’s most valuable firms today are those pouring resources into AI, racing one another to build ever-greater machines. Whether you watch the news, sit in a corporate boardroom, or glance at a university’s latest policies, AI is everywhere. AI chip companies are worth more than countries.<sup>41</sup> AI spending has ‘added more to GDP growth than consumer spending’.<sup>42</sup> The gold rush Altman envisioned is well underway.

Alas, we are left to reckon with the other half of his prediction – the part about *ending the world*. We have got our great companies. Now, I suppose, we must prepare for the end of the world.

---

39 Nick Lichtenberg, ‘AI Has Added \$160 Billion to “True GDP” Since 2022, Goldman Sachs Says. There’s Just One Problem: That’s Not Official’, *Fortune* (Web Page, 17 September 2025) <<https://fortune.com/2025/09/17/how-much-gdp-artificial-intelligence-goldman-sachs-160-billion/>>.

40 nova clave, ‘Fireside Chat: Sam Altman President, YCombinator and Mike Curtis, VP of Engineering, Airbnb – 2015’ (YouTube, 21 February 2024) 00:08:47–00:08:56 <<https://www.youtube.com/watch?v=d6IDZpvHAoo&t=527s>>.

41 See, eg, Derek Saul, ‘How Valuable Is Nvidia? More Than 97% Of Countries’ Economies – And Much More’, *Forbes* (Web Page, 9 July 2025) <<https://www.forbes.com/sites/dereksaul/2025/07/09/whats-nvidia-more-valuable-than-the-uk-economy-worlds-30-richest-and-total-nyc-real-estate-make-the-list/>>. See also Noel Randewich and Shashwat Chauhan, ‘Nvidia Briefly on Track to Become World’s Most Valuable Company Ever’, *Reuters* (Web Page, 4 July 2025) <<https://www.reuters.com/business/nvidia-set-become-worlds-most-valuable-company-history-2025-07-03/>>.

42 Randy Watts, ‘Capital Spending as the Key Market Driver?’, *Forbes* (Web Page, 14 August 2025) <<https://www.forbes.com/sites/randywatts/2025/08/14/capital-spending-as-the-key-market-driver/>>, quoting Neil Dutta, Head of Economic Research at Renaissance Macro Research.