and Stockholm conventions in 2019 following a 15-year debate. Obviously, implementing this qualified majority voting would be politically challenging and feasible only if voting was used sparingly and transparently, and without overlooking or dismissing the concerns of countries with less power.

Seize the moment

The choices made now will determine whether the health of people and the planet are safeguarded or put at further risk. We urge the INC's newly elected chair to consider implementing these reforms.

Multiple events during the past few years have undermined multilateralism. And as geopolitical priorities shift, environmental concerns are increasingly being sidelined or environmental policies weakened[10].

Against this backdrop, it is crucial that the INC process succeeds — both to address a major contributor to the interconnected planetary crises (climate change, biodiversity loss and pollution) and to restore faith in the idea that international cooperation can solve global challenges.

The authors

Paul Einhäupl is a research associate of the Systemic Risks Group at the Research Institute for Sustainability (RIFS) at GFZ Helmholtz Centre for Geosciences, Potsdam, Germany. **Linda Del Savio** is a research associate of the Ocean Governance Group at the Research Institute for Sustainability (RIFS) at GFZ Helmholtz Centre for Geosciences, Potsdam, Germany. **Melanie Bergmann** is a senior scientist at the Alfred Wegener Institute, Helmholtz Centre for Polar and Marine Research, Bremerhaven, Germany. **Annika Jahnke** is co-head of the Department of Exposure Science at the Helmholtz Centre for Environmental Research — UFZ, Leipzig, and professor of exposome analytics at RWTH Aachen University, Aachen, Germany.
e-mail: paul.einhaeupl@rifs-potsdam.de

1. Bergmann, M. *et al. Nature Rev. Earth Environ.* **3**, 323–337 (2022).
2. Tekman, M. B., Walther, B. A., Peter, C., Gutow, L. & Bergmann, M. *Impacts of Plastic Pollution in the Oceans on Marine Species, Biodiversity and Ecosystem* (WWF, 2022).
3. Stoett, P. *One Earth* **5**, 593–596 (2022).
4. Karali, N., Khanna, N. & Nihar, S. *Climate Impact of Primary Plastic Production* (Lawrence Berkeley National Laboratory, 2024).
5. Dauvergne, P. *et al. Mar. Policy* **181**, 106820 (2025).
6. Maes, T., Wienrich, N., Weiand, L. & Cowan, E. *Camb. Prisms Plast.* **1**, e22 (2023).
7. MacLeod, M., Arp, H. P. H., Tekman, M. B. & Jahnke, A. *Science* **373**, 61–64 (2021).
8. UNEP. 5/14. *End Plastic Pollution: Towards an International Legally Binding Instrument: Resolution/Adopted by the United Nations Environment Assembly* (UN, 2022).
9. Monclús, L. *et al. Nature* **643**, 349–355 (2025).
10. Luo, H. & Sun, Y. *Sci. Rep.* **14**, 10747 (2024).

The authors declare no competing interests.

Does AI already have human-level intelligence? The evidence is clear

Eddy Keming Chen, Mikhail Belkin, Leon Bergen & David Danks

By any reasonable criteria, the vision of human-level machine intelligence laid out by Alan Turing in 1950 is now a reality. Eyes unclouded by dread or hype will help to prepare for what comes next.

In 1950, in a paper entitled 'Computing Machinery and Intelligence'[1], Alan Turing proposed his 'imitation game'. Now known as the Turing test, it addressed a question that seemed purely hypothetical: could machines display the kind of flexible, general cognitive competence that is characteristic of human thought, such that they could pass themselves off as humans to unaware humans?

Three-quarters of a century later, the answer looks like 'yes'. In March 2025, the large language model (LLM) GPT-4.5, developed by OpenAI in San Francisco, California, was judged by humans in a Turing test to be human 73% of the time — more often than actual humans were[2]. Moreover, readers even preferred literary texts generated by LLMs over those written by human experts[3].

This is far from all. LLMs have achieved gold-medal performance at the International Mathematical Olympiad, collaborated with leading mathematicians to prove theorems[4], generated scientific hypotheses that have been validated in experiments[5], solved problems from PhD exams, assisted professional programmers in writing code, composed poetry and much more — including chatting 24/7 with hundreds of millions of people around the world. In other words, LLMs have shown many signs of the sort of broad, flexible cognitive competence that was Turing's focus — what we now call 'general intelligence', although Turing did not use the term.

Yet many experts baulk at saying that current AI models display artificial general intelligence (AGI) — and some doubt that they ever will. A March 2025 survey by the Association for the Advancement of Artificial Intelligence in Washington DC found that 76% of leading researchers thought that scaling up current AI approaches would be 'unlikely' or 'very unlikely' to yield AGI (see go.nature.com/4smn16b).

What explains this disconnect? We suggest that the problem is part conceptual, because definitions of AGI are ambiguous and inconsistent; part emotional, because AGI raises fear of displacement and disruption; and part practical, as the term is entangled with commercial interests that can distort assessments. Precisely because AGI dominates public discourse, it is worth engaging with the concept in a more detached way: as a question about intelligence, rather than a pressing concern about social upheaval or an ever-postponed milestone in a business contract.

In writing this Comment, we approached this question from different perspectives — philosophy, machine learning, linguistics and cognitive science — and reached a consensus after extensive discussion. In what follows, we set out why we think that, once you clear away certain confusions, and strive to make fair comparisons and avoid anthropocentric biases, the conclusion is straightforward: by reasonable standards, including Turing's own, we have artificial systems that are generally intelligent. The long-standing problem of creating AGI has been solved. Recognizing this fact matters — for policy, for risk and for understanding the nature of mind and even the world itself.

Questions of definition

We assume, as we think Turing would have done, that humans have general intelligence. Some think that general intelligence does not exist at all, even in humans. Although this view is coherent and philosophically interesting, we set it aside here as being too disconnected from most AI discourse. But having made this assumption, how should we characterize general intelligence?

A common informal definition of general intelligence, and the starting point of our discussions, is a system that can do almost all cognitive tasks that a human can do[6,7].

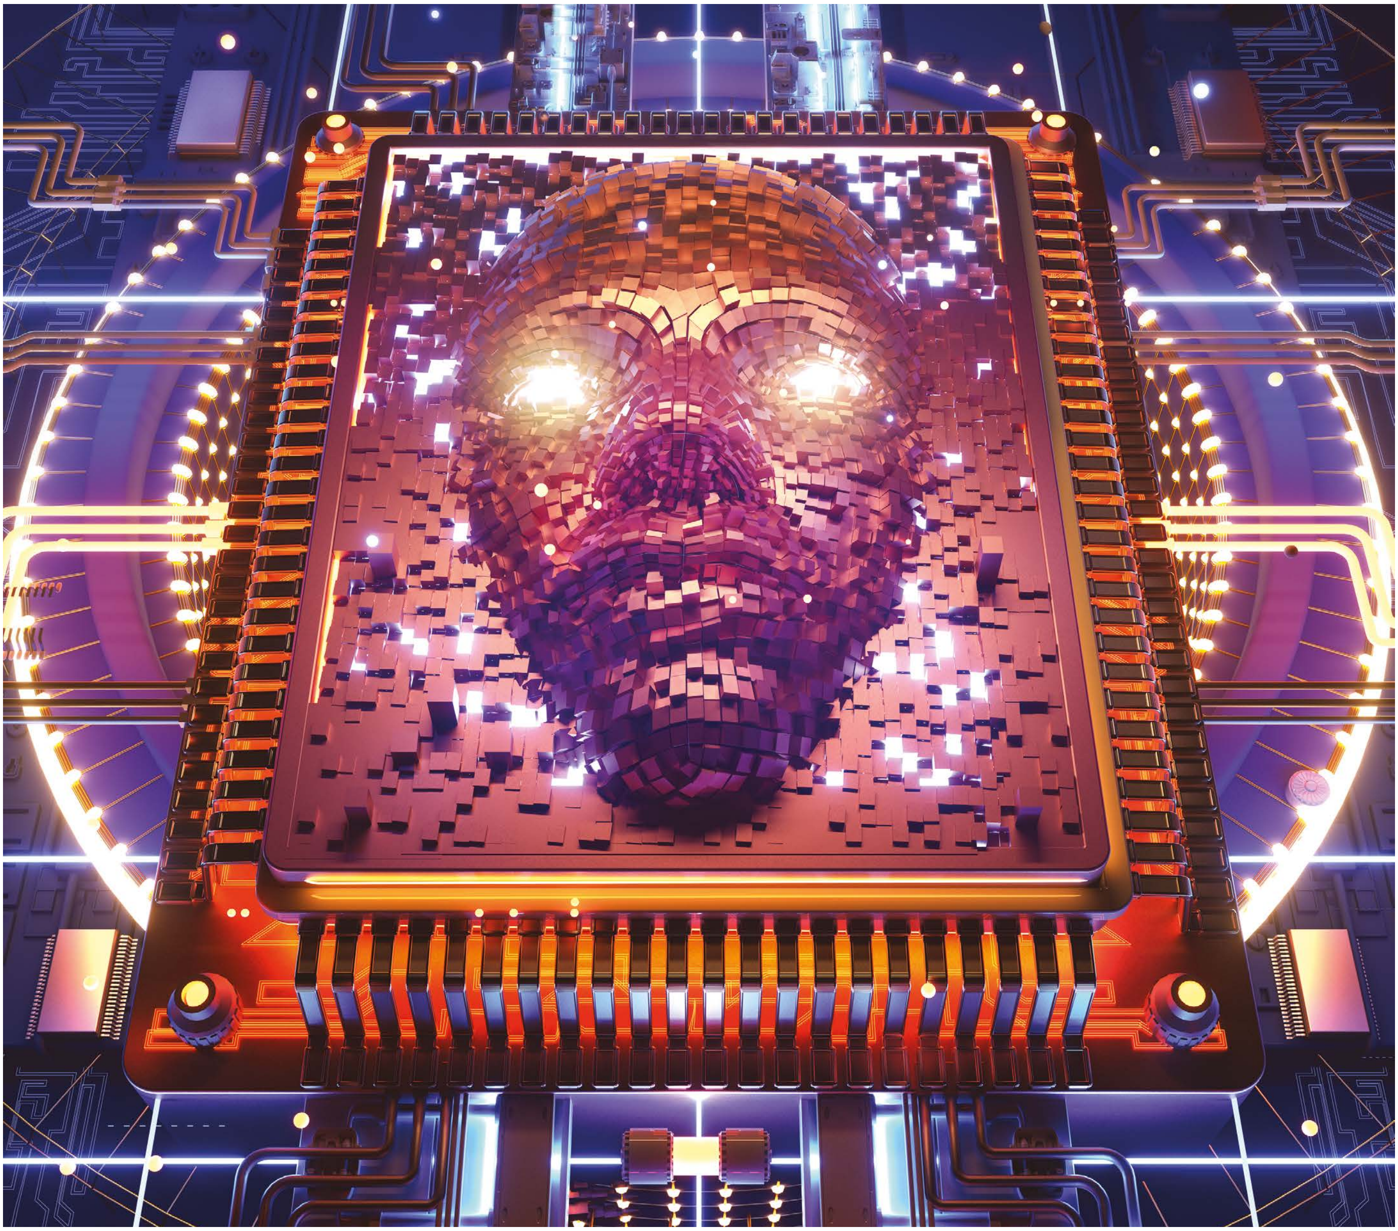

ILLUSTRATION BY JACEY

What tasks should be on that list engenders a lot of debate, but the phrase 'a human' also conceals a crucial ambiguity. Does it mean a top human expert for each task? Then no individual qualifies – Marie Curie won Nobel prizes in chemistry and physics but was not an expert in number theory. Does it mean a composite human with competence across the board? This, too, seems a high bar – Albert Einstein revolutionized physics, but he couldn't speak Mandarin.

A definition that excludes essentially all humans is not a definition of general intelligence; it is about something else, perhaps ideal expertise or collective intelligence. Rather, general intelligence is about having sufficient breadth and depth of cognitive abilities, with 'sufficient' anchored by paradigm cases. Breadth means abilities across multiple domains – mathematics, language, science, practical reasoning, creative tasks – in contrast to 'narrow' intelligences, such as a calculator or a chess-playing program. Depth means strong performance within those domains, not merely superficial engagement.

Human general intelligence admits degrees and variation. Children, average adults and an acknowledged genius such as Einstein all have general intelligence of varying level and profile. Individual humans excel or fall short in different domains. The same flexibility should apply to artificial systems: we should ask whether they have the core cognitive abilities at levels comparable to human-level general intelligence.

"Insofar as individual humans have general intelligence, current LLMs do, too."

Rather than stipulating a definition, we draw on both actual and hypothetical cases of general intelligence – from Einstein to aliens to oracles – to triangulate the contours of the concept and refine it more systematically. Our conclusion: insofar as individual humans have general intelligence, current LLMs do, too.

What general intelligence isn't

We can start by identifying four features that are not required for general intelligence.

Perfection. We don't expect a physicist to match Einstein's insights, or a biologist to replicate Charles Darwin's breakthroughs. Few, if any, humans have perfect depth even within specialist areas of competence. Human general intelligence does not require perfection; neither should AGI.

Universality. No individual human can do every cognitive task, and other species have abilities that exceed our own: an octopus can control its eight arms independently; many insects can see parts of the electromagnetic spectrum that are invisible to humans. General intelligence does not require universal mastery of these skills; an AGI does not need perfect breadth.

Human similarity. Intelligence is a functional property that can be realized in different substrates — a point Turing embraced in 1950 by setting aside human biology[1]. Systems demonstrating general intelligence need not replicate human cognitive architecture or understand human cultural references. We would not demand these things of intelligent aliens; the same applies to machines.

Superintelligence. This is generally used to indicate any system that greatly exceeds the cognitive performance of humans in almost all areas. Superintelligence and AGI are often conflated, particularly in business contexts, in which 'superintelligence' often signals economic disruption. No human meets this standard; it should not be a requirement for AGI, either.

A cascade of evidence

What, then, is general intelligence? There is no 'bright line' test for its presence — any exact threshold is inevitably arbitrary. This might frustrate those who want exact criteria, but the vagueness is a feature, not a bug. Concepts such as 'life' and 'health' resist sharp definition yet remain useful; we recognize paradigm cases without needing exact boundaries. Humans are paradigm examples of general intelligence; a pocket calculator lacks it, despite superhuman ability at calculations.

When we assess general intelligence or ability in other humans, we do not attempt to peer inside their heads to verify understanding — we infer it from behaviour, conversation and problem-solving. No single test is definitive, but evidence accumulates. The same applies to artificial systems.

Just as we assess human general intelligence through progressively demanding tests, from basic literacy to PhD examinations, we can consider a cascade of increasingly demanding evidence that warrants progressively higher confidence in the presence of AGI.

Turing-test level. Markers comparable to a basic school education: passing standard school exams, holding adequate conversations and performing simple reasoning. A decade ago, meeting these might have been widely accepted as sufficiently strong evidence for AGI.

Expert level. Here, the demands escalate: gold-medal performance at international competitions, solving problems on PhD exams across multiple fields, writing and debugging complex code, fluency in dozens of languages, useful frontier research assistance as well as competent creative and practical problem-solving, from essay writing to trip planning. These achievements exceed many depictions of AGI in science fiction. The sentient supercomputer HAL 9000, from director Stanley Kubrick's 1968 film *2001: A Space Odyssey*, exhibited less breadth than current LLMs do. And current LLMs even exceed what we demand of humans: we credit individual people with general intelligence on the basis of much weaker evidence.

Superhuman level. Revolutionary scientific discoveries and consistent superiority over leading human experts across a range of domains. Such evidence would surely allow no reasonable debate about the presence of general intelligence in a machine — but it is not required evidence for its presence, because no human shows this.

Turing's vision realized

Current LLMs already cover the first two levels. As LLMs tackle progressively more difficult problems, alternative explanations for their capabilities — for instance, that they are gigantic 'lookup tables'[8] that retrieve pre-computed answers or 'stochastic parrots'[9] that regurgitate shallow regularities without grasping meaning or structure — become increasingly disconfirmed.

Often, however, such claims just reappear with different predictions. Hypotheses that retreat before each new success, always predicting failure just beyond current achievements, are not compelling scientific theories, but a dogmatic commitment to perpetual scepticism.

We think the current evidence is clear. By inference to the best explanation — the same reasoning we use in attributing general intelligence to other people — we are observing AGI of a high degree. Machines such as those envisioned by Turing have arrived. Similar arguments have been made before[10] (see also go.nature.com/49p6voq), and have engendered controversy and push-back. Our argument benefits from substantial advances and extra time. As of early 2026, the case for AGI is considerably more clear-cut.

We now examine ten common objections to the idea that current LLMs display general intelligence. Several of them echo objections that Turing himself considered in 1950. Each, we suggest, either conflates general intelligence with non-essential aspects of intelligence or applies standards that individual humans fail to meet.

They're just parrots. The stochastic parrot objection says that LLMs merely interpolate training data. They can only recombine patterns they've encountered, so they must fail on genuinely new problems, or 'out-of-distribution generalization'. This echoes 'Lady Lovelace's Objection', inspired by Ada Lovelace's 1843 remark and formulated by Turing as the claim that machines can "never do anything really new"[1]. Early LLMs certainly made mistakes on problems requiring

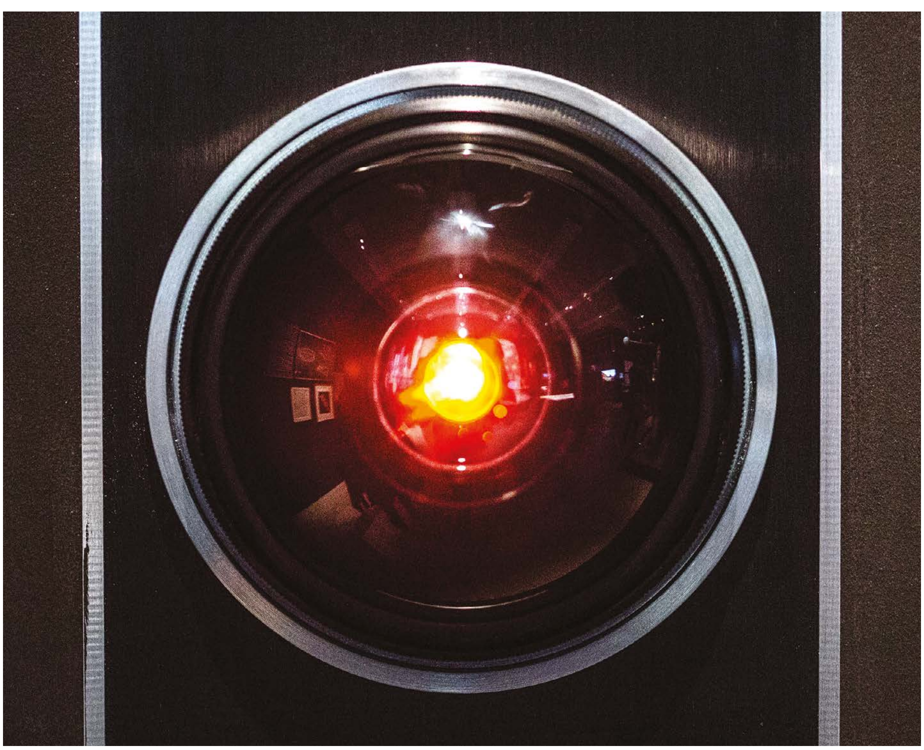

Current AIs are more broadly capable than the science-fiction supercomputer HAL 9000 was.

HETHERS/SHUTTERSTOCK

reasoning and generalization beyond surface patterns in training data. But current LLMs can solve new, unpublished maths problems, perform near-optimal in-context statistical inference on scientific data[11] and exhibit cross-domain transfer, in that training on code improves general reasoning across non-coding domains[12]. If critics demand revolutionary discoveries such as Einstein's relativity, they are setting the bar too high, because very few humans make such discoveries either. Furthermore, there is no guarantee that human intelligence is not itself a sophisticated version of a stochastic parrot. All intelligence, human or artificial, must extract structure from correlational data; the question is how deep the extraction goes.

They lack world models. LLMs supposedly lack representations of their physical environment that are necessary for genuine understanding. But having a world model requires only the ability to predict what would happen if circumstances differed — to answer counterfactual questions. Ask a cutting-edge LLM what differs between dropping a glass or a pillow on a tile floor, and it will correctly predict shattering in one case and not the other. The ability of LLMs to solve olympiad mathematics and physics problems and assist with engineering design suggests that they possess functional models of physical principles. By these standards, LLMs already have world models. Furthermore, neural networks developed for specialized domains such as autonomous driving are already learning predictive models of physical scenes that support counterfactual reasoning and sophisticated physical awareness[13].

They understand only words. This objection centres on the fact that LLMs are trained only on text, and so must be fundamentally limited to text-based tasks. Frontier models are now trained on images and other multimodal data, making this objection somewhat obsolete. Moreover, language is humanity's most powerful tool for compressing and capturing knowledge about reality. LLMs can extract this compressed knowledge and apply it to distinctly non-linguistic tasks: helping researchers to design experiments — for example, suggesting what to test next in biology and materials science[4] — goes beyond merely linguistic performance. We are yet to encounter the sharp limitations to LLM performance that this objection predicts.

They don't have bodies. Without embodiment, critics argue, there can be no general intelligence. This reflects an anthropocentric bias that seems to be wielded only against AI. People would ascribe intelligence to a disembodied alien communicating by radio, or to a brain sustained in a vat. An entity that

GERARD NOONAN/ALAMY

Alan Turing asked whether machines could think.

responds accurately to any question, but never moves or acts physically, would be regarded as profoundly intelligent. Physicist Stephen Hawking interacted with the world almost entirely through text and synthesized speech, yet his physical limitations in no way diminished his intelligence. Motor capabilities are separable from general intelligence.

They lack agency. It is true that present-day LLMs do not form independent goals or initiate action unprompted, as humans do. Even 'agentic' AI systems — such as frontier coding agents — typically act only when a user triggers a task, even if they can then automatically draft features and fix bugs. But intelligence does not require autonomy. Like the Oracle of Delphi — understood as a system that produces accurate answers only when queried — current LLMs need not initiate goals to count as intelligent. Humans typically have both general intelligence and autonomy, but we should not thereby conclude that one requires the other. Autonomy matters for moral responsibility, but it is not constitutive of intelligence.

> **"A difference in efficiency of learning does not necessarily mean a different level of intelligence."**

They don't have a sense of self. Critics object that AGI requires persistent autobiographical memory, stable personal identity and continual self-updating, in a way that today's LLMs lack. Frontier LLMs are increasingly deployed with long-term context and user-specific memory, so the gap in memory and identity between humans and LLM-based systems is shrinking. Although these features are key to how humans function in society, they are not requirements for general intelligence. We do not deny intelligence to humans with profound amnesia, who might remember almost no personal details, nor to people with multiple personality disorder, whose personalities might be unaware of each other. For example, an intelligent amnesiac could reason powerfully using information recorded in external notebooks, even if those notebooks never connect.

They are inefficient learners. A common objection is that children learn concepts from few examples, whereas LLMs require vast amounts of data: they have low 'sample efficiency'. Even if true, the comparison is not straightforward. It ignores billions of years of evolutionary 'pre-training' that built in rich inductive biases — about objects, space and causation — long before learning from experience begins. More importantly, a difference in efficiency of learning does not necessarily mean a different level of intelligence. A chess master who takes ten years to achieve mastery is just as good a chess player as someone who reaches that level in one year.

They hallucinate. LLMs sometimes confidently present false information as being true, raising concerns about reliability.

Hallucination is becoming less prevalent in current models, but it does not in any case disqualify general intelligence. Humans are prone to false memories, cognitive biases and perceptual illusions, often asserting them confidently.

They lack economic benefits. In some circles, particularly in industry, the definition of AGI has moved to mean something that must generate substantial economic returns. But economic ability is an application of general intelligence, not a requirement for its presence. History offers many examples of brilliant individuals who generated little economic return; we do not thereby deny them general intelligence.

Their intelligence is 'alien'. LLMs sometimes fail at tasks that are trivial for us (such as counting the number of instances of the letter 'r' in 'strawberry') while performing well on tasks humans find difficult – from PhD-level science questions to rapid synthesis across vast scientific literature collections. But our claim is not that LLMs have human intelligence; it is that they realize some form of general intelligence. Moreover, frontier LLMs increasingly write code and use tools to compensate for their weaknesses, just as humans augment their abilities with technology, from pocket calculators to smartphones. The resulting profile of strengths and weaknesses is a fairly alien form of intelligence, but that is a reason to broaden our conception of general intelligence, not to deny that these systems have it.

Throughout these objections, critics demand particular characteristics that no solid principle requires of general intelligence. Many exclude intelligences we readily recognize as such; some of them exclude humans considered exceptionally intelligent, or all humans altogether.

Finally, a different reaction, not an objection as such, is what Turing called "heads in the sand": the consequences of machines thinking would be too dreadful, so let us hope they cannot. This is understandable from the emotional and human point of view. But as Turing noted, it calls for consolation, not refutation – one can sympathize with the worry without treating it as an argument.

Why this matters

Recognizing current LLMs as AGI and as fulfilling the vision of machine intelligence set out by Turing is a wake-up call. These systems are not on the horizon; they are here. Frameworks designed to assess narrow tools are inadequate for evaluating their benefits and risks. Questions of coexistence, responsibility, liability and governance take on new dimensions when the systems involved are not narrow instruments but general intelligence.

Intelligence, as we have seen, does not require strong autonomy – a finding that complicates debates about legal and moral responsibility of artificial-intelligence systems, which often assume the two go together. We need more careful, empirically grounded ways to assess and establish responsibility for AI. Furthermore, conventional methods of governance are unlikely to work for AGI, precisely because of its generality. Technology is typically governed on the basis of its possible uses, but AGI can be used almost anywhere.

Another key question concerns the relationship between human intelligence and the forms of general intelligence that have been artificially created and will be created in the future. In many ways, these systems are surprisingly human-like – they write like us, talk like us and share some of our imperfections. Yet they remain alien, reflecting a fundamentally different path to general intelligence, unconstrained by the evolutionary pressures that shaped the survival-driven goals of human cognition, a small squishy body, scarce energy and low-bandwidth communication. Understanding this alienness matters. How do these systems differ from ours, and are those differences transient or fundamental?

> **"For the first time in human history, we are no longer alone in the space of general intelligence."**

The answers could reveal which aspects of general intelligence are universal and which are parochial features of our biological inheritance. For the first time in human history, we are no longer alone in the space of general intelligence. We might also better understand the risks, because an alien general intelligence might fail in surprising ways, or succeed in ways that cannot be easily understood or guided. Seeing these systems for what they are will help us to work with them today and prepare for what comes next.

Just five years ago, we didn't have AGI; now we do. Even more powerful forms of intelligence will no doubt arrive soon. This is both remarkable and concerning. Remarkable, because we are privileged to witness what is perhaps the most significant scientific and technological revolution in human history. Concerning, because the timeline is compressed beyond any historical precedent and could be accelerating.

In 1965, reviewing AI progress for the RAND Corporation, philosopher Hubert Dreyfus compared approaches to developing human-level artificial intelligence to trying to reach the Moon by climbing a tree (see go.nature.com/3ywerhj). For decades, this seemed right. But as the evidence mounts, it is increasingly clear we misjudged the nature of the Moon and the power of the tree. General intelligence can indeed emerge from simple learning rules applied at scale to patterns latent in human language – patterns rich enough, it turns out, to encode much of the structure of reality itself.

Nicolaus Copernicus displaced humans from the centre of the cosmos. Darwin displaced humans from a privileged place in nature. Turing suggested that humans might not embody the only way to be intelligent. The machines Turing envisioned 75 years ago have finally arrived, in a form both more alien and more human than anyone imagined. Like those earlier revolutions, this one invites us to rethink our standing – and accept that there are more kinds of minds than we had previously entertained. Our place in the world, and our understanding of mind, will not be the same.

The authors

Eddy Keming Chen is an associate professor of philosophy at the University of California, San Diego, San Diego, California, USA. **Mikhail Belkin** is a professor of artificial intelligence, data science, and computer science at the University of California, San Diego, San Diego, California, USA. **Leon Bergen** is an associate professor of linguistics and computer science at the University of California, San Diego, USA, and a member of technical staff at Goodfire AI in San Francisco, California, USA. **David Danks** is a professor of data science, philosophy and policy at the University of California, San Diego, San Diego, California, USA.
e-mail: eddykemingchen@ucsd.edu

1. Turing, A. M. *Mind* **LIX**, 433–460 (1950).
2. Jones, C. R. & Bergen, B. K. Preprint at arXiv https://doi.org/10.48550/arXiv.2503.23674 (2025).
3. Chakrabarty, T., Ginsburg, J. C. & Dhillon, P. Preprint at arXiv https://doi.org/10.48550/arXiv.2510.13939 (2025).
4. Bubeck, S. *et al.* Preprint at arXiv https://doi.org/10.48550/arXiv.2511.16072 (2025).
5. Rizvi, S. A. *et al.* Preprint at bioRxiv https://doi.org/10.1101/2025.04.14.648850 (2025).
6. Morris, M. R. *et al. Proc. 41st Int. Conf. Mach. Learn.* **235**, 36308–36321 (PMLR, 2024).
7. Hendrycks, D. *et al.* Preprint at arXiv https://doi.org/10.48550/arXiv.2510.18212 (2025).
8. Block, N. *Phil. Rev.* **90**, 5–43 (1981).
9. Bender, E. M., Gebru, T., McMillan-Major, A. & Shmitchell, S. in *Proc. 2021 ACM Conf. Fairness Account. Transparency* 610–623 (Association for Computing Machinery, 2021).
10. Bubeck, S. *et al.* Preprint at arXiv https://doi.org/10.48550/arXiv.2303.12712 (2023).
11. Dai, Y., Gao, Z., Sattar, Y., Dean, S. & Sun, J. Preprint at arXiv https://doi.org/10.48550/arXiv.2506.07298 (2025).
12. Ma, Y. *et al.* Preprint at arXiv https://doi.org/10.48550/arXiv.2309.16298 (2023).
13. NVIDIA. Preprint at arXiv https://doi.org/10.48550/arXiv.2501.03575 (2025).

The authors declare competing interests; see go.nature.com/4bvblgb for details.
Supplementary information accompanies this article online; see go.nature.com/4bvblgb for details.