

A user's guide to the end of the world (as we know it)

What if we knew that the world — the human world — would be so radically different within our lifetimes that we might not recognize daily life? What if we knew that children born in 2025 would never know the meaning of work, or income inequality, or deprivation?

What if the ensuing shocks were so profound — to society, business, government, even to our sense of self — that our future selves wished more than anything else they had prepared better for the day when algorithms and machines could do everything we do, only better and faster? And what if our future selves were to look back at 2024, to see it as the one clear moment in time when we saw the future and blinked?

The potential for those shocks is there. Artificial General Intelligence — software with human-like intelligence and the ability to self-teach — may be nearing a state where it can, at least theoretically, start to displace, at scale, the functions (mental, physical, perhaps even emotional) that have, for millennia, made humans the species we are.

Will the resulting shocks come in a decade or a century, or somewhere in between?

In the long arc of time, the timing may not matter, as we know today the clock is running out on the age in which man reigned over machine. We are on the edge of a new era of commingling and interoperability, an era which could see intelligent machines play a role in every aspect of life.

That era will present plenty of unknowns, and for that, society needs to start preparing.

To discuss how, a group of technologists, academics and executives gathered this spring in Asilomar, California, to confront our newest existential challenge: What if we succeed? What if, within a decade, AGI is capable of replicating every human task? And how, on earth, should we prepare?

The Setting

Past is prologue, and so may be true for AI in the quiet solitudes of Asilomar.

Jutting into the Pacific Ocean, around the corner from Monterey Bay, Asilomar is a sleepy retreat that can easily be bypassed for the spectacle of Pebble Beach on one side and chill vibe of Pacific Grove on the other. Indeed, it seems to embody the paradox that Ezra Pound applied to her hometown of Oakland, not far away. There is no 'there' there.

In one direction lies the sea and its infinite promise, and in the other, beyond the coastal mountains, Silicon Valley and its exponential promise. It was here, at the edge of America, and in the throws of the second Industrial Revolution, that the great San Francisco architect Julia Morgan designed a retreat for the Young Women's Christian Association, the first in the American West.

Morgan had already helped San Francisco rebuild from the Great Fire of 1906, and was well into the defining commission of her career, Hearst Castle, down the coast in San Simeon. In 1913, when Asilomar opened, the world was on the cusp of a technological revolution, one that would make airplanes, cars, telephones and movies the tent pegs of 20th century life.

Morgan had tried to give the YWCA a retreat from what was to come, with her wood-beam and vaulted ceiling "Chapel" and its engraved words, "the lord on high

is mighty." But she later recognized that the Roaring Twenties, and the rise of American modernism, would challenge that view of the almighty, as it gave god-like powers to a new scientific class immersed in the atom and electron.

Nearly a century later, a new generation of scientists, technologists and their backers, are seeking to equally reshape society, with AGI.

Can they prepare for the unknowns of a far more powerful tech revolution? Can they find ways for autonomous machines and their human dependents to co-habitat?

If they succeed, can their society reform capitalism, in ways that the first Roaring Twenties failed to do, to fairly distribute resources even though the means of production are controlled by machines? And will the rest of us find new ways to accept our finiteness in an infinite economy?

Looking over the rugged dunes that connect Asilomar with the setting sun, and the promise of tomorrow, the challenges of technology disruption may feel the same today as they did for Julia Morgan, to both harness modernity and keep it in its place. And yet a century on, this new revolution feels entirely different, with its exponential promises looking to be as profound as its existential threats.

Here are some of the considerations:

1. The promise of infinite surplus

Moore's Law remains the guiding force of our times, allowing for the doubling of computing power every two years. In fact, over the past decade, compute power has doubled every six months. To date, human ingenuity has been able to keep pace with that kind of growth. We've figured out how to use computers, smart phones, and wired machines to our benefit. But the compounding of compute beyond this decade, into a new AGI realm, may be more challenging to human adaptation, especially as machines increasingly make their own decisions and gain physical mobility. We could see more biorobots in 25 years than humans, and far more virtual agents informing, advising, and eventually directing our daily lives.

Some forecasters believe that 80 per cent of jobs will be done by some form of AI within a generation. One result could be a near-infinite rise in economic output, as the world's productive capacity soars. Services such as health care, education and financial advice could become free, universally available and ever-improving. Profound challenges — climate, cancer, crime — could be solved rapidly. And even as wages collapse with the end of work, spiking productivity rates and surging output should easily compensate, if effective distribution models are established.

While the promise of such surpluses may be foreseeable, the timing is not. Like all technological impacts, AI is following the course of a slow, steady explosion. That makes it harder for society to prepare, to change income models, tax regimes, and social expectations. Moreover, the course of AI adoption and its impacts are unlikely to be linear, especially when they run up against rigid social and economic models. A meandering path to AGI, with technological bursts and social reactions, may mean we get to AGI before society is ready for it.

Perhaps Transformational AI can help distribute the surpluses it creates, but only if we guide it to do both at the same speed.

What's needed:

Dynamic research to track the displacement of labour and distribution of benefits across the economy.

2. What AI needs to learn

Many aspects of AI may not be as smart as we think, given the stumbles of ChatGPT. But it's also showing all signs of being slow and steady, then fast and furious. The avalanche effect.

The technology is currently evolving through a rapid series of small steps, with few eureka moments. One reason: most models still focus on computation, rather than achieving goals. The chat bot explosion of the early 2020s has yet to expose deep thinking from machines, other than an impressive ability to accept prompts and respond.

We need to shift Large Language Models (LLMs) to algorithms that can learn from ordinary experiences, not just from neat data sets. Indeed, LLMs may need to search for greater challenges, and pursue the sorts of messes that literally don't compute. AI may also need more time, tools, and space to test and learn from multiple hypotheses rather than the hard coding of symbolic reasoning. Ultimately, systems will need to get better at working with the wonders of the human mind, and the depths of intuition that machines can't replicate. One suggestion: "Think of it more as parenting than programming."

That kind of parental guidance — helping AI not touch the hot stove or poke the dog's eye — can come through continuous deep learning and "efficient off-policy learning"; that is, allowing the models to colour outside the boundaries of their algorithms, to understand aberrations, and to engage in discordant and shallow data sets. This will require humans to accept that our relationship with AI increasingly will become continuous and not transactional; again, parenting, not daycare. And we may need to accept that passive systems will wake up to learning.

Ultimately, AI will need to develop its ability to anticipate and adjust. Prediction machines will need to become planning machines. And like most of us, there will need to be plans for failure. Hospitals, for instance, may need to add on-call data scientists to help manage algorithms that go awry or stop during a procedure because they don't know what to do.

The growth of AI may be iterative, a journey of baby steps. But such is the rapidly incremental nature of innovation.

A bit like childhood.

What's needed:

Open sharing of discoveries and data, when public interest is at stake, to ensure collective progress.

3. A concentration risk

You don't need to be in Silicon Valley to hear the giant sucking sound of capital by AI. And it's getting louder, as the colossi of chip, cloud, and compute devour more and more capital to finance their energy- and data-hungry learning models. As the big get bigger, they're also starting to drive returns, which in turn is leading to more capital generation.

LLM spending is already estimated to have hit about \$1 billion last year, and could reach \$10 billion this year or next. Some suggested \$100 billion could be spent annually on language models within five years. Intel is already spending \$25 billion on chips. AI is having the same power in fundraising; last year saw \$50 billion in venture funding and 38 new unicorns. OpenAI, the market darling, saw its valuation edge reach \$80 billion.

And then there's this calculation: If AGI increases economic productivity, in an optimistic forecast, the cash value of its benefits could be \$124 quadrillion. Suddenly, a \$7 trillion investment seems reasonable.

The centripetal force of AI is about more than money. The cloud behemoths are accumulating data and talent at rapid clips, and also amassing the resources to spend on supercomputers. It's estimated fewer than 10,000 people are working on what can be considered "transformative" AI — anything that might lead to AGI — and most are serving the interests of a handful of firms. It's said that Tesla's dominance in automobile data, especially for autonomous vehicles, was one reason Apple — hardly a constrained enterprise — backed away from the its AV project.

Will the concentration lead to an oligopoly or even monopoly in AI? And will that stifle competition? Or will there be an emergence of "bilateral oligopolies" — small groups of players at each link in the supply chain? That could lead to cartels or at least coalitions in, for instance, electricity supply, computing operations, and chip supplies. Governments could equally impose constraints — quotas, as an example — on dominant players, or at least require them to serve national needs first.

All of which comes with a caution: concentration of power is less dangerous than concentration of thought.

What's needed:

Governments may need to consider an industrial policy mix for AI, to ensure a fair and strategic allocation of resources, including capital.

4. A risk to supply chains

There may be more than we can manage. Compute, chips, and labour are all in short supply, and traditional supply-demand models may no longer apply. For one, AI is creating exponential curves in demand through the unpredictability of its uses and needs. The steep cost of inferencing — the running of data in a live AI model — is only growing as those models get hungrier. The more they learn, the more they want to learn. A separate tech race is on, to develop more efficient chips, shrink the size of models, and compress the middleware that adds more weight to systems. In each of those areas, competition helps. And the explosion of capital for AI could help fuel that competition.

Structural (or infrastructural) inputs like electricity will be harder to fix. Much of the world is already in a hurry to produce more clean electricity to run factories and cities in a net-zero economy, and there's a risk that capital-rich AI projects and their energy-hungry data centres will outbid the older parts of the economy trying to transition their energy models. In that scenario, the compute demands of AI could sideline the climate demands of society. In those cases, governments may need to assign scarce resources to a hierarchy of societal needs.

More positively, the enormous potential of the race for AGI, and the apparent economic potential, could prove to be an added incentive to the development and scaling of emerging energy sources like nuclear fusion.

A scarcity of inputs will also challenge the business and organizational adoption of AI, including transformational AI. Legacy industries, already operating with low margins, will continue to be challenged to compete, compute, and to buy the chips they may need. Such a scenario may lead many companies and public-sector organizations, as was the case in the Internet's early years, to accept their place as slow adopters, using off-the-shelf enterprise software tools that can be useful for efficiency but less dynamic for innovation.

This will put further pressure on governments, to find ways to increase both supply and demand for AI in a broad range of sectors as well as public interest pursuits. As is often said of the Internet, we had an invention that was profound and powerful enough to cure cancer, and we used it instead to share photos. The same risk — individual preferences versus collective needs — could play out with AI and models; creating celebrity avatars rather than diagnosing health problems. In business, too, the next generation of AI needs to be focussed on discovery, not just automation. Collective research models, such as a DARPA or NASA for AI, could help coordinate university research and business application, and in turn develop ecosystems that ease supply chain constraints and open doors for emerging challengers.

Ultimately, AI should expand our vision, not shrink it.

What's needed:

Incentives and initiatives to ensure the supply chains of AGI are focussed on societal needs, especially science, including the incomplete sciences of climate and behaviour.

5. A risk to robots

Mention AI on Main Street, and most conversations will quickly turn to robots and their rise. The early years of Transformational AI is painting a different picture. Many of the biggest private sector AI players have set aside their initial focus on blue-collar work — where robots are most needed — and turned instead to white-collar functions. For one, there's quicker returns in the information economy. By its very nature, language models are also best at playing with words and numbers, the stuff of enterprise software. And it turns out, error rates are more acceptable in the information economy. We're willing to accept fake news, or fake essays, a lot more than flawed buildings.

That's not to suggest there's no hope for robots outside warehouses. It's just going to take longer. Big Tech is actively trying to develop software that can mimic human dexterity and senses. The prize is enormous. It just takes an ability to convert perception data into action data — what we might call reflex and instinct, as opposed to habit. In the coming years, we may see more “teleportation,” as people take possession of robots to help them learn. We could even see business models around Brain as a Service, in which enterprise software packages can be bought or licensed to command various aspects of the workplace, home, and community, or perhaps even ourselves.

The demand for robots, and other smart hardware, will only grow as populations age and eventually shrink. So, too, will our comfort interacting with machines, just as we're comfortable conversing with our phones. (One retailer said their store tests show customers trust on-floor robots more than on-floor staff, for information.)

What will AI-powered robots, and other learning machines, be good for? If we get it wrong, we'll end up developing self-teaching vacuum cleaners and toilet scrubbers first, rather than using Transformational AI to transform how the world's economy operates. If we get it right, AGI can help remove transportation from the ground and sea, putting it in the air and freeing up our lands and waters for better uses. It can transform manufacturing, including through 3D printing. And most profoundly, it can change the way we live, with medical devices in our bodies learning as we age. Like the third Industrial Revolution — the computer age — which allowed us to shift en masse from a brawn economy to a brain economy, the advancement of Transformational AI can power the robots and smart machines in our lives to do more than make our lives more convenient and efficient.

They can help us leap into a new age of discovery.

What's needed:

Robotics programs, including public supports, that drive innovation to the most important frontiers of human progress.

6. A transition risk

Utopia doesn't have an on-ramp. If we're to get to an AI-driven world, in which there's infinite surpluses and machine-enabled peace and prosperity, we will have to endure a lot of bumpy detours and diversions.

In the world's poorest countries, and indeed in the poorest regions of the world's richest countries, labour is too abundant and cheap to replace with AI. Infrastructure and technology distribution will further impede the universal spread of AI. Paradoxically, where AI is needed most, it could be deployed least.

The dispersion of AI in advanced economies won't come without disruptions, either, especially to workforces. Entire areas of expertise, and the trades and professions associated with them, could rapidly dwindle, along with the education programs that feed them. "Stranded expertise," as it's called.

During this transition, many of us will need to shift to "augmented work" in which we job-share with AI, exploring ways to make the most of each other as we co-habitate roles. We will also need to prepare — psychologically as well as economically — for the day when we're no longer needed in that role. Augmentation will give way to an advanced form of automation, in which the job and its constituent tasks continue to evolve in the hands of a machine.

Those with a growth mindset see far more opportunity. First of all, if AI is restricted to current human knowledge, it will have failed. Properly guided, Transformational AI should multiply our collective knowledge set, as well as our troves of creativity, which in turn will lead to more discoveries, more creations and more pursuits and jobs. As one small comparison, the microscope did not eliminate any jobs; rather, it opened our collective eyes to frontiers and possibilities we had scarcely imagined.

Bumps, yes, but the transition is to a place of greater human engagement.

What's needed:

Development programs for AI in low-income regions, as well as AI-powered learning programs across professions, trades, and jobs at risk.

7. A distribution risk

Even if we put AI in Utopia, it will be subject to human nature, which generally is not about sacrifice and sharing. Yes, once AGI becomes a universal reality, the potential surpluses of our economy could spell an end to hunger, poverty, and disease. But humans may not be content. We may still need and yearn for status hierarchies. Our happiness will remain relative. There will also be divisions between countries, as nations (xenophobic ones, especially) seek forms of differentiation to enhance national pride and self-worth. An AI-powered Olympics would be no fun if the optimal outcome was for every country to share the gold medal.

This kind of competition — or as Freud called it, “the narcissism of small differences” — may become more entrenched, and violent, if humans are unable to find other forms of meaning, beyond work. Regardless of the political economy of a country, basic instincts will be a challenge for AI to cope with — something communist states discovered about themselves and their Utopian dreams in George Orwell’s *Animal Farm*. (“All animals are equal, but some animals are more equal than others.”)

Even today, in the West at least, we have the best lives humanity has arguably ever lived, and yet we generally feel we don’t have enough. Social discontent has rarely been higher, and ironically, we know how to solve most of society’s shortcomings. In fact, we don’t need AI to figure out how to distribute wealth more equitably, as we did that some generations ago. Just open our borders more to trade and immigration, and find more systematic ways to distribute the surpluses of our economies. AI would tell us to do the same thing, presumably, and we would find reasons — relative prosperity — to reject it.

What’s needed:

More open trade policies, including for digital assets and IP, to allow for a freer flow of AI opportunities and benefits.

8. A risk to democratic capitalism

Capitalism exists by permission of democracy, and if the benefits of AI are not clearly and fairly distributed, the system that is financing its growth could be at risk. This could require capitalism to adjust as much as society needs to adjust to the powers of AI.

For centuries, the distribution of economic surpluses has been largely based on labour. More recently, economic rewards have gone disproportionately to the owners of capital, over labour. As AI, and the owners of the capital behind it, amass more economic benefits, and as labour rewards are diminished, social tensions and ensuing political pressures could grow. This could become even more acute in aging societies in which older, and less productive, generations hold the bulk of capital through their lifetime of savings, while labour-challenged younger generations are squeezed.

Could this lead governments to nationalize AI, in order to distribute the benefits more widely? Or will governments instead more aggressively tax the owners of capital, to redistribute their gains from AI? Perhaps modern capitalism won’t be needed anyway, since its AGI may replace the need for markets to determine equilibriums and drive the efficient allocation of resources. An algorithm can do that.

As AGI takes hold, governments could also be tempted by policies more associated with authoritarianism, to maintain control over the social and political consequences of emerging models. Fundamentally, democratic capitalism will be challenged to address this: Whoever controls the digital infrastructure behind AI — supercomputers, chips, energy sources — will control the future. In other words, the digital means of distribution will eclipse the means of production as the determinant of economic power.

Which leads to this question: in 2034, if Silicon Valley hasn't taken over Washington, will Washington need to take over Silicon Valley?

What's needed:

Businesses, investors, and governments need to rapidly develop new approaches to market economics, to ensure the rewards of capital and labour are properly assessed and allocated.

9. A risk to meaning

Technology has always challenged the meaning of life, and the purpose we each hold. Deus ex machina (“god from the machine”) goes back to Ancient Greece, and a seemingly instinctive association between the almighty and technology, both being stronger than us. In ancient theatre, the god from the machine usually brought resolution to the problems on centre stage and sent audiences home happy. AGI may be expected to do the same, even though the angst of human life may not compute.

Humans will need to prepare, perhaps rapidly, for a world in which work and deprivation are both remarkably scarce. That won't put an end to human desires, even when everyone has sufficient food, housing, and clothing. We always need more. Especially in our minds and hearts. AGI may not anytime soon be able to speak to our emotional needs, for laughter, comfort, and love. Nor can it address the social isolation that can come from the end of workplaces, schools, and commercial centres.

Or can it?

AGI may actually not put an end to work, but rather enhance jobs and pursuits with more meaning. It will take the robotic out of every job, perhaps. This could lead to a new definition of work, in which jobs are as much social as economic functions. Call it a Seinfeldian world, as someone suggested, each of us busy with banter and errands. We'll all be active, and rewarded accordingly, just not what exactly what we're sure for. Will that shift to leisurely work make us feel more inconsequential? And perhaps less essential? Will it lead to lethargy? Or anarchy?

Over the coming years and decades, as we pursue the final frontiers of technology, we will need to explore the inner frontiers of humanity, to determine what it means to be humans. We can love and preserve, as much as we today produce and provide. But that will require some new shared narratives of what the good life — and good work — can be.

Only humans can code that.

What's needed:

Dismantle or at least refine labour market barriers and regulations, to allow for a more entrepreneurial, creative, and human approach to work.

10. A risk to regulation

The greatest risk in regulation may be our inclination to regulate the past against the future, and AGI is all about the future. **That presents an important moment to challenge ourselves with what ifs:**

- What if there is only one AI model and it can be independently regulated?
- What if we regulate the users and not the algorithms?
- What if we declare and code all models with what good looks like? What if we declare and code all models with what bad looks like, including self-replication, break-ins and evil intent (e.g. bioweapon design)?
- What if we ensure agents and models have “normative competence” to search for, and recognize, boundaries and laws?
- What if we penalize, even threaten to shut down, models that go against good?
- What if we use interoperability to monitor how models are doing, and ultimately allow models to measure and police themselves?
- What if we allow models to share IP, to assist new entrants?
- What if we require AI models that draw on data from public spaces — roads, social channels, education systems, for instance — to join data utilities?
- What if we create regulatory safe harbours for areas of public importance, such as disease recognition?
- What if we assign “personhood,” with rights and legal responsibilities, to agents and chatbots?
- What if we apply principles rather than prescriptions to AI?
- And ultimately, what if beneficial co-existence is not possible?

The emerging frontiers of AI regulation are no longer in the distance, and governments (democratic ones, at least) will be challenged to catch up. Fearing the worst, they may throw in the towel and shut down AGI efforts — or leave it in the hands of incumbent oligopolies that may be easier to negotiate with and police. It's surely the case that AGI is too novel a concept to allow for regulatory capture. And yet, the incumbents, and their regulators, are party to the rise of algorithms that may soon be too complex and inscrutable for them to understand, and dangerously irreversible.

There is no easy way at it, other than, perhaps, to remind ourselves that science is inherently about experiment, guided by universal principles, including Do No Harm. Societies, in a range of political systems, have harnessed the benefits of science — space, medical, nuclear, biological — by following such principles. Ultimately, we may need to place the same confidence in the scientists working on AGI. If we don't, other countries and regimes will not likely let up in their pursuit of this new frontier for intelligence. We may be better to work together, and over time, as was the case in the atomic age, place faith in science and a bit of skepticism in each other.

As the political code suggests, trust but verify.

What's needed:

In the near term, a clear and replicable taxonomy and code for AI regulators to model and share. In the longer term, international conventions and systems for AI governance.

11. A risk to global security

Scientists hate to be politicized. Too late. AI is rapidly becoming a central political issue, and a growing geopolitical one. The G7 is making AI one of its top priorities, in part to ensure there's a coherent and collective approach to keep China and Russia from achieving supremacy. The United States and Britain have made AI a central file for their heads of government, as nuclear security was in decades past. They're not alone. The United Arab Emirates, among other emerging economic powers, has made AI a national ambition, while its close ally India is seeking to do the same with what may be the fastest growing tech stack anywhere. Those challengers to the West may find their own common ground, in a "Third Way" model that is neither Chinese, nor American-centric.

A space race in AI may be healthy for competition, and innovation, but it's also a risk to global security, as self-learning models strive to compete with each other based on national standards and goals, not universal ones. This rivalrous approach to AI could deepen as countries put more resources behind national strategies designed to create a competitive advantage. Potentially worse may be national restrictions (and hoarding) of key AI inputs, including compute power and chips. Without greater global governance, the odds of mishaps — intentional or accidental — will grow.

Fortunately, the world has nearly a century of experience in successful multilateral governance, which while flawed, has helped prevent nuclear strikes, the proliferation of biological weapons, and ultimately another world war. Even conventions on child labour, land mines, and summary executions have had their effect. Similar approaches to AI governance may soon be needed.

Unfortunately, the post-war institutions that have successfully governed conduct in so many areas since the 1940s are themselves under attack. If the major powers are losing confidence in the World Trade Organization, why would they lean into a World AI Organization? As in previous generations, it may be up to scientists and business leaders to build bridges with all countries pursuing AI goals, including those that may have difficult political relationships with others. As the Churchillian credo of diplomacy says, jaw jaw is better than war war. In that spirit, we will need more alignment, between East, West, North, and South, on the goals — and dangers — of AI. We will also need more public confidence in AI, for people to see the value in its development as well as global governance, understanding its weaponization would be fatal.

Ultimately, AI for all will require all for AI.

What's needed:

Track 2 diplomacy to bring together scientists, business leaders and academics from rival countries, paving the way for a Track 1.5 effort with government officials.

12. A risk to society

Not far from the barren dunes and windswept groves of Asilomar, the great midcentury American writer John Steinbeck worked on *The Grapes of Wrath* and *Cannery Row*. Those classics captured America at a crossroads, scarred by Depression, challenged by a changing world order, and yet inspired by the technological gusto from the Roaring Twenties. Writing of an emergent superpower, Steinbeck noted that the best qualities that Americans seek in people — kindness, honesty, openness — are not what they value in the market. And what we seek in markets — sharpness, acquisitiveness, self-interest — are what we consider failures in people. In other words, we seek in a system what we don't want in each other, failing to appreciate a system is a function of its parts.

Can AI change that, taking the best of humanity and applying it to the worst of society? It won't be easy given the dyspeptic mood of publics almost anywhere. It will be even harder in a political environment that seems to eschew kindness and celebrate sharpness.

The mind-boggling reach of Transformational AI can seem like too much for any society to comprehend and absorb. Democracy, most of all, may be challenged to mediate those existential challenges. The risks to our personal and collective security, the dangers of concentration, the unknowns of distribution, and the highly variable outcomes of regulation — each of these could tip the public's mind away from AI. That is, if Transformational AI is not too fantastical for the public to consider seriously. That is, if it's not too late to reverse what's been started. That is, if we can untangle what's smarter, faster, and more aware than its creators.

And if we can, do we know how to move collectively and at speed? As a society, we weren't ready for the COVID-19 pandemic, which was predictable and precedented. Facing the unprecedented, we will need to find a different path. We can start by breaking down challenges into actionable and meaningful opportunities, and to frame the AGI discussion in the realities of today and tomorrow, rather than the extraordinary projections of a future time. Governments and their publics care most about the here and now, which is a good place to meet. Taking a page from nuclear science, we can also develop the muscles and rigours of safety precautions and monitoring. And we can build bridges with scores of countries to ensure this is a human-scale endeavour, not the purview of an elite band. Steinbeck wrote, in *Cannery Row*, "Man's right to kill himself is inviolable, but sometimes a friend can make it unnecessary." That may sound morbid, but it was framed in the spirit of a community that was overwhelmed by the changing world around it. Friendship, they discovered, was one of humanity's great powers.

It may yet be what prepares us for the age of AGI.