Foreword

Joshua Achiam



The Apollo 11 Saturn V rolls out of the Vehicle Assembly Building
Kennedy Space Center, May 20, 1969. NASA.

AN INESCAPABLE AND TIMELESS FACET OF OUR RELATIONSHIP WITH NATURE IS conflict: man versus nature. The natural environment is not hospitable to humans; it is filled with chances for death, disease, struggle, and want. However, we have three great advantages in our long contest with nature. The first is statecraft, which allows us to band together and cooperate to reduce collective risk. The second is our capacity for invention, which allows us to modify the natural environment to better support us. The third is the fact that we adapt to nature much faster than nature adapts to us. Nature is not a totally static and unchanging adversary, but on the whole we can think and solve problems faster than it can evolve new problems for us.

By nature, I mean something more expansive than just the plants and animals and weather that make up the environment. I mean everything that exists inside the universe but outside human will. I include the physical frailty of our bodies, germs and diseases, the behaviors of cells and their tendency to mutate into cancers. I include the complexity of the laws of physics and how much is hidden from us by virtue of operating at a length or timescale we can't perceive. I mean almost everything in the physical world, and the complete set of affordances or constraints that nature imposes on our lives.

The project of statecraft as we practice it today is organized, directly or indirectly, around our struggle with nature. Certainly, the principle that "if men were angels, no

government would be necessary" accounts for much of what states are required to do—a great deal of the practical work is about managing the relationships of people with each other and with institutions—but the deeper structural motivations come back to nature, to what is or isn't possible under the laws of physics, to the limits of the human body and mind.

If nature did not impose scarcity of resources, we would not need to have contests for them, we would not need markets to organize their distribution, and we would not need force to defend their trade. If nature did not contain the possibility of pain, we could not be coerced or deterred by it. If we did not have limitations on our intelligence, wisdom, or observational capacity, we could reliably find the positive-sum cooperative solutions to our disputes. If we did not need each other's cooperation to survive and if we could not impose costs on each other, perhaps it would not matter whether men were angels.

AGI

A strange thing is happening: our capacity for invention is leading to the imminent development of artificial general intelligence (AGI), defined roughly as AI systems that are as capable as the most adept humans at most things. As I write this foreword, the frontier AI models that exist today are already nearly as capable as the best human experts in many domains; in some, far beyond them. There is an increasingly firm consensus in the AI research community that AGI is within reach.

The arrival of AGI will be transformative. It evokes comparisons to the discovery of agriculture, the Industrial Revolution, and the information age. Another common analogy, less wonderful but certainly containing its own dark majesty, is to nuclear weapons. These analogies are helpful, but each is incomplete, as no prior transformation of the world entailed so many rapid and simultaneous potential disruptions. Every analogy from the past played out over decades or centuries, but it is expected that AGI's arrival will be marked by extraordinary speed, with seismic shifts playing out over single-digit years. The speed is owed to a mix of factors: first, software is much easier to distribute than other technologies; and second, by virtue of its nature, AGI will help unpack itself and put itself to work.

Many of the most important problems in the world are bottlenecked by intelligence and by the number of experts who know how to attempt any attack at all, let alone solve them. This doesn't mean just the highest-hanging fruit of grand challenges like string theory: think of every problem in engineering design, materials synthesis, manufacturing processes, law, logistics, strategy, and so on. In every field of study and practice, human expertise is scarce and correspondingly expensive, and progress is rate-limited by the availability of expertise. An AGI that is able to perform at or above human expert level would massively change the cost structure of applying expertise: the cost of a solution will be the lower of the machine cost or the human cost, and the machine cost will drop over time as compute becomes more abundant and models become more efficient.

This changes the arrival rate of progress. We have built our states and institutions around expectations of a certain number of breakthroughs per decade in crop yields, military technology, energy production, and other research-driven fields. For example, we built the global system of energy production—including physical infrastructure, trade relationships, force posture, and diplomacy—largely around assumptions that alternative energy sources to oil would be slow to develop technologically. But AGI that automates research and development stands to radically change these types of assumptions, and correspondingly many objectives of states will have to change.

Similarly, we have built bureaucracies through which the state provides services under assumptions of human bottlenecks in information sharing, process development, and decision-making that may soon be obviated by the application of AGI. The scope of what can be provided for the common good at affordable cost could massively expand. What we stand to unlock in knowledge, to gain in medical science, to accelerate in agriculture, to unleash in energy, and to provide at raw materials cost will unburden us from many of the limits from nature that have historically shaped our design of government.

There is another profound difference between the arrival of AGI and any previous technological revolution in history: this technology has, to some extent, a mind of its own. It is plausible that AGIs will turn out to have goals that are not identical to ours. The science of alignment is still nascent, and there is not a guarantee that we will produce perfectly obedient machines. Even if most of them are mostly obedient, we may well be adding something to the environment that is beyond our will. If nature includes AGIs that have minds of their own, we can no longer count on the persistence of our third advantage—our ability to outspeed nature in a contest. This will pose a new and foundational challenge for statecraft to address.

Futurecraft

The essays that follow take the arrival of AGI as a premise and ask what we can do in the early phase of transforming government, industry, and society. The authors offer ideas on how we might marshal our efforts to achieve the best possible outcome for the American people.

The questions they address are simultaneously far-seeing and grounded. How can an increasingly automated executive branch remain accountable? Who bears responsibility for AI-agent swarms acting in the economy? What electricity networks, industrial capacity, and scientific infrastructure will be required to put abundant intelligence to work? The proposals range from new administrative records and legal identities for AI agents to parallel power networks and a simulated human body. Across these domains, the authors follow the implications of AGI down to the records, responsibilities, and physical systems on which civilization depends.

The work of redesigning the state in response to AGI will be an immense task made difficult by the compressed time frame. The machinery of government is slow, but the need and demand for action will be extraordinary. Politicians today are only just wak-

ing up to the existence of AGI; the full set of transformations it will bring is still difficult to talk about without sounding a little bit crazy. But the future does not require us to believe in it beforehand for it to arrive all the same. The speed of events means that every scrap of preparation undertaken in advance will help.

The institutions and practices established in the early years may become the assumptions everyone else builds on. Every effort made now to map the frontier of statecraft could have an outsized impact, perhaps even a strong founder effect, on what comes next. I feel that energy—the potential for precedent—in these essays.

Perhaps, though, the best way to describe this collection is to think of the prompt it would take to ask a future super intelligence to generate it. If I had to guess, it would be something like: "Build the future of American statecraft for the day after AGI. Set humanity up for the best possible future."

"Make no mistakes."

ABOUT THE AUTHOR

Joshua Achiam is a non-resident senior fellow at the Foundation for American Innovation currently focused on how rapid progress in AI may create new public policy and national security needs. He previously worked in AI safety research and leadership at OpenAI for nine years, serving as Head of Mission Alignment in 2025 and as Chief Futurist in 2026. He holds a PhD in electrical engineering and computer science from UC Berkeley and bachelor of science degrees in aerospace engineering and physics from the University of Florida.



From *Futurecraft: The Day After AGI*, Foundation for American Innovation, 2026.
Read or listen online at futurecraft.thefai.org/foreword.

© 2026 by the Foundation for American Innovation. All rights reserved.