To: Faisal D'Souza, NCO
Office of Science and Technology Policy
2415 Eisenhower Avenue
Alexandria, VA 22314

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

Maintaining U.S. AI Dominance via Prioritizing Semiconductors, Energy, & Security

Ross Gruetzemacher, Assistant Professor of Business Analytics
Pierre Harter, Associate Vice President for Research and Industry Engagement

Introduction[1]

This response to the Request for Information (RFI) for the U.S. AI Action Plan focuses on four independent components driving U.S. AI dominance: semiconductors, energy, security, and algorithms. Semiconductors, energy, and security are things which are easier to maintain a leadership position in because dominance is correlated with investment in existing, well-established technologies. Algorithms are trickier because progress is driven by investments in human capital; there is a greater risk that we will fall behind China — e.g., simply consider the sheer number of Chinese STEM graduates relative to the U.S. — and there are grave national security risks that arise from racing China on algorithmic progress.

Here, we argue that by racing ahead in semiconductors, energy, and security, while also implementing America First policies designed to constrain the efforts of both our adversaries and our international partners, we could achieve a position of such strength that if we acted responsibly, we would be insulated from the potentially catastrophic risks of an unrestrained AI algorithmic race. DeepSeek R1's surprising capabilities were perhaps a Sputnik moment for the U.S., but the biggest takeaway from this should be that the algorithmic race is going to remain close. Regardless of who leads in AI algorithms, the lead will always be slim; if we lead in semiconductors, energy, and security, we will still dominate the future of AI even if we fall behind on algorithms.

[1] This proposal is inspired by a unique perspective that combines 1) our institution's broad and unbiased view across the industry and geopolitical landscape and 2) six years of experience designing and facilitating tabletop exercises simulating AI races between states and leading AI firms.

Our Proposed Strategy

Succinctly, we propose a strategy of winning at all costs on semiconductors, energy, and security, while being cautious and restrained with respect to algorithmic innovation. For racing ahead in semiconductors and energy we recommend large investments, tax incentives, and deregulation. For racing ahead with respect to security we propose a novel, holistic approach to AI security that we call AI Security, Evaluation, and Control (AI-SEC). AI-SEC addresses not only cyber offensive threats but also national security threats such as from AI accidents or misuse, and relies on regulation and, if necessary, the partial nationalization of AI firms.

Leading — and ultimately winning — in semiconductors, energy, and security will also yield long-term advantages in reaping the economic benefits of AI. To understand this, we distinguish between leading in algorithmic development and leading in algorithmic deployment; those who have a small lead in algorithms will see limited economic benefits if they have limited capacity for deployment because the future deployment of advanced AI systems will be closely coupled to semiconductors, energy, and security.

Moreover, years of experience wargaming AI arms races suggests that we must reach international agreements regarding algorithmic development to ensure that all states pursuing advanced AGI are able to do so responsibly and avoid the two primary categories of risk:

1. Risks inherent in the technology itself such as catastrophic AI accidents or misuse
2. Risks from conflict arising between racing nations to sabotage adversaries' algorithmic development efforts[2]

If the U.S. maintains a lead in algorithmic deployment, we will retain a much stronger position from which to negotiate with China or other racing nations regarding agreements on responsible and cautious algorithmic development to reduce category one risks. While the credible threat of a deterrence regime like nuclear mutually assured destruction (e.g., '*MAIM*'; Hendrycks et al. 2025) may also be necessary for reducing category two risks, we must avoid acting on that threat at all costs — leading in algorithmic deployment can be the carrot to the threat of destruction stick.

[2] Such as extreme cyber attacks on energy infrastructure or kinetic attacks to destroy AI computing infrastructure. Attacks like this could have much further reaching catastrophic impacts to civilians, either by default or if they led to escalating retaliatory strikes.

Our recommendations are likely to stand out from others responding to this RFI (e.g., Anthropic) in one prominent way: we are urging cautious and restrained algorithmic development to mitigate what we see as the two most concerning classes of national security risks. Our proposed approach is necessary because previous arms races only concerned risk from an adversary whereas this race involves a second class of risk inherent in the technology itself. If the inherent risk is taken seriously, we cannot expect strategies from previous arms races to be able to be effective in isolation. Moreover, while there are strategic advantages that go to the first actor to develop any technology, arguments for AI winner-take-all scenarios are tenuous — unlike the uncontrolled nuclear chain reaction we were racing toward in the Manhattan Project we don't even have an agreed up definition for what would constitute AGI, advanced AGI, machine superintelligence, or any other terms that could bring about radical power shifts and societal transformation that is often discussed (e.g., '*decisive strategic advantage*'; Bostrom 2014).

While comparisons with the Manhattan Project may not be apt for all elements of the new AI Action Plan, a World War II-style production effort (Herman 2013) to enhance cybersecurity for AI systems and other systems deemed critical to national security (e.g., energy and semiconductor production infrastructure) could mitigate many of the risks. Moreover, if the administration is especially concerned with the speculative strategic risks of radically transformative AI technologies, a Manhattan Project-style algorithmic development effort would be more responsible than continuing to allow a half dozen or more private firms to race to such a powerful technology. If the tenuous arguments for strategic risks hold water, the U.S. is at risk of being undermined by whichever AI firm reaches whatever it is that constitutes radically transformative AI first (e.g., 'decisive strategic advantage'; Bostrom 2014).

While a Manhattan Project-style project would be safer than unrestrained development in the private sector, neither of these paths are ideal for the current circumstances. Rather than complete nationalization in a war-like big science project — e.g., the Manhattan Project or the Apollo Program — a public-private partnership in this era, particularly with the need to reward investors and AI firms for their work, would be more appropriate. Specifically, we feel that heavy regulatory steps while preparing the framework for a web of public-private partnerships, unprecedented in complexity, is prudent. To maximize innovation, the President would need to identify a risk threshold after which the public-private partnerships would begin to be implemented, perhaps in a staged manner, such that when the technology approaches nuclear weapons level risks it is firmly managed in partnership between a designated federal agency, the Department of Defense, and an AI laboratory or a broader consortium of AI laboratories and possibly experts from national labs and academia.

Experience from AI wargames suggests that public-private partnerships are very effective means of managing catastrophic AI risks that may arise from races between AI laboratories. Regardless of what is decided on the topic of public-private partnerships and nationalization, the U.S. needs to thoroughly consider all options, and consider conducting its own wargames to better understand the complex issues involved in AI races.

Four Components to AI Dominance

There are four critical components to successfully maintaining U.S. dominance in the Artificial Intelligence (AI) race: semiconductors, energy, security, and algorithms.[3] It is critical that the U.S. continues to race forward with the first three of these so that we're able to ensure secure algorithmic development and optimize deployment. We describe these four components and what the government's role might look like in the following pages.[4]

Semiconductor & Energy Dominance

AI chips are the most important of these components with respect to the economic impacts of AI, and this is the case even before we get to advanced AGI. Consider that, as AI continues to become increasingly capable, performing more and more economically valuable tasks, AI chips will become a means of production able to supplant white collar labor. There has been some debate over the impact that generative AI will have on the labor market, and while economists disagree about the

[3] We note that the first three of the four components—semiconductors, energy, and security—significantly overlap with the core elements of Anthropic's response to this RFI. However, concerningly, algorithmic progress is conspicuously missing from Antrhopic's response. AI firms have a self-interest in deregulation, but certain technologies should not be developed by public firms. Consider that the U.S. has the greatest free market in the world, and we can boast of the greatest defense sector of any nation, but nuclear weapons are not developed, produced, or maintained by the defense sector; nuclear weapons are developed, produced, and maintained by the Department of Energy in partnership with the Department of Defense. No nuclear state trusts nuclear weapons development to private firms. If AI is going to have as much or more strategic significance as nuclear weapons, then it would be a very bold decision to allow its development, production, and maintenance to be managed entirely by the public sector. Nationalization, or a full-scale war-effort big science project like the Manhattan Project or the Apollo Program may not be necessary, but some form of public-private partnership should be considered, and plans should be prepared to act soon, if necessary. Again, this is something that is seen quite frequently in wargames, and it is a more successful strategy than the alternative.

[4] We also note that this proposal is inspired by four years of experience running a tabletop exercise simulating an AI race with China. This experience is described in great detail in the paper Strategic Insights from Simulation Gaming of AI Race Dynamics (Gruetzemacher et al. 2025). The most significant takeaway is that there is no winner when the race between the U.S. and China is not resolved through agreements to ensure the safe and responsible development of advanced AGI.

severity of AI’s impact, there is widespread agreement that AI will automate many of the tasks that humans previously performed (see, e.g., Autor 2024).

Whether or not the labor that AI displaces will reskill and find new employment is not relevant here; what is relevant is that AI will soon displace a significant portion of cognitive work in the U.S. and abroad.[5] This labor will be replaced by semiconductors using large amounts of electricity housed in data centers. The only limit to the economic value that AI will be able to generate will be access to the energy and the AI chips that power these data centers. AI chips will effectively expand labor markets throughout the globe, and U.S. economic leadership will hinge on controlling or profiting from as much of this newly created silicon labor force as possible. Dario Amodei, CEO of Anthropic, suggests that upon the arrival of advanced AGI we will begin to see instances of a “country of geniuses in a datacenter”. Therefore, it is essential that the U.S. controls and profits from the great wealth that will be generated from our semiconductor technologies.

Racing to develop domestic semiconductor production and to rapidly expand domestic energy production to power domestic data centers is critical to U.S. near- and long-term economic prosperity and national security. The value of pursuing U.S. semiconductor and energy supremacy is independent of the value of a high-risk pursuit of domineering leadership in AI algorithms. China is not far behind in semiconductor design, and it will be difficult to prevent Chinese firms from accessing the computational resources necessary to train competitive if not the leading models. However, rigorously enforced export controls over all levels of the semiconductor supply chain could decisively prevent China and Chinese firms from being anywhere close to the level of the U.S. and U.S. firms in the theater of AI deployment. In general, all policy levers should be considered to increase AI chip production as quickly and as much as possible, including more extreme options such as invoking the defense production act.

Energy is nearly as critical to AI deployment dominance as the semiconductors themselves. This is due to the fact that, unless American citizens are ready to pay for the cost of an AI race, we need to rapidly expand our energy infrastructure to be prepared to support the energy demands of large-scale AI deployment. If we fall behind here, Americans’ electricity prices are likely to increase dramatically as AI begins to displace labor. The U.S. should use all tools at its disposal to address this critical issue of national security, including embracing both fossil fuels and renewables. Massive government spending and an unprecedented effort to expand our energy production capacity are necessary. AI is going to increase productivity, and will drive radically

[5] If readers have not yet seen what OpenAI’s ChatGPT Pro version of Deep Researcher can do, we encourage them to try to witness this capability firsthand. This is the most convincing evidence to date of where AI systems are quickly heading.

growth to GDP, so spending on energy, critical to national security, which will help drive the explosive economic growth that generative AI will lead to, can be justified.

Nuclear energy has some of the greatest potential, but, if we wanted to rely on this we would need to cut all regulatory red tape, and place our trust in 80 years of U.S. industry efforts to develop safe nuclear power. The U.S. has one of the world's best records for nuclear safety, and no expense should be spared for safety, but we should not let regulatory burden slow the process. Because leaning into nuclear power is prudent for our proposed effort, we should consider dramatic steps such as recommissioning any decommissioned reactors, even if just for a short amount of time, and we must act on this as soon as possible.

Security & Algorithms

There is a tremendous amount at stake with the development of advanced AGI. We have eluded to this earlier, but below we provide some prominent quotes.

"*Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last, unless we learn how to avoid the risks*."
– Stephen Hawking, 2014

"... the leader in this sphere will become the ruler of the world."
– Vladmir Putin, 2018

"*With artificial intelligence, we are summoning the demon. You know all those stories where there's the guy with the pentagram and the holy water and he's like... yeah, he's sure he can control the demon, [but] it doesn't work out.*"
– Elon Musk, 2014

Musk is not alone among tech leaders. Other prominent leaders from Silicon Valley are not secretive about their ongoing efforts to build bunker compounds to retreat to in the event of a global catastrophe or civilization collapse (e.g., Mark Zuckerberg and Peter Thiel). While Superforecasters are skeptical of human extinction, they have placed the risk of permanent disempowerment of humanity by advanced AGI at five percent.[6]

Four years of TTX wargame experience provides evidence that there are many different ways that AI could precipitate catastrophe, including destabilization to the point of civilizational collapse or a reshuffling of the global order as a result of the AI race

[6] https://goodjudgment.com/superforecasting-ai/

leading to a conventional world war. One of the most prominent themes that emerged from four years of simulating AI races was the frequency of extreme cyberwarfare. Semiconductor technologies, data centers, and algorithmic innovation are all likely targets of potential cyber offensive efforts not only from China, but from all of the usual suspects.

The strategy we propose here is not granular or comprehensive, and it is not ready for implementation as policy. Whatever strategy the President selects to include in the AI Action Plan should not follow any single proposal or response provided from the responses to the RFI. While not entirely compatible, we feel that what we propose in this response could complement other proposals well (for example, Hendrycks et al. 2025). We will not go into detail about such possibilities here due to length restrictions but are open to further engagement.

To address the vast array of national security risks that AI poses, including both the inherent risks of the technology and the risk from an uncontrolled escalation of a great power conflict, we propose a holistic approach to security that we refer to as AI Security, Evaluation, and Control (AI-SEC). AI-SEC is intended to balance the security concerns from three independent types of risks that could result in catastrophic consequences for humanity:

- Potentially uncontrolled escalation of a great power conflict (wartime risks)
- Risks inherent to AI systems (i.e., accident risks)
- Threats from irresponsible use by misguided actors (misuse risks)

These three types of risks are just a more granular representation of the two categories of risks discussed earlier, with the first category being broken down further into accident and misuse risks. It is easy to see how these three types of risks map to the three components of AI-SEC. Security maps to mitigating risks associated with traditional national security concerns like great power conflict; evaluation maps to preventing misuse risk; control maps to preventing accident risk. Only with a holistic approach like AI-SEC can we be confident that we are being responsible in our efforts to mitigate the broad spectrum of catastrophic national security risks that the development of advanced AI presents.

Conclusion

Our key recommendations are summarily listed below:

- We recommend exercising all policy levers to expedite the rapid development of semiconductor production capacity in the U.S.

- We recommend reducing all regulatory oversight and taking emergency executive action to accelerate the construction and recommissioning of energy production facilities, including all forms of energy production, such as fossil fuels, nuclear, and renewables.
- We recommend public-private partnerships between the U.S. and AI laboratories to ensure that strategic technology on par with nuclear weapons is managed responsibly
    - As part of the process to develop an AI Action Plan, we encourage the President to commission a series of wargames, combining AI expertise from labs and academia with leaders from national security agencies, to better understand AI race dynamics.
- We propose a new, holistic regime of security for the race to radically transformative AI systems that we call AI Security, Evaluation, and Control (AI-SEC).
    - AI-SEC balances the national security concerns from risks associated with a potentially uncontrolled escalation of a great power conflict (wartime risks) with threats from risks inherent to AI systems (i.e., accident risks) with threats from irresponsible use by misguided actors (misuse risks).
        - Security addresses traditional national security risks, such as wartime risks.
        - Evaluation addresses threats from misuse of AI.
        - Control addresses threats from losing control over AI systems.

References

Anthropic, 2025. Re: Request for Information (RFI) on the Development of an Artificial Intelligence (AI) Action Plan ("Plan"). Submitted by D'Souza, F. in response to 2025-02305 (90 FR 9088).

Bostrom, N., 2014. Superintelligence. Oxford University Press.

Autor, D., 2024. Applying AI to rebuild middle class jobs (No. w32140). National Bureau of Economic Research.

Gruetzemacher, R., Avin, S., Fox, J. and Saeri, A.K., 2025. Strategic Insights from Simulation Gaming of AI Race Dynamics. Futures, p.103563.

Hendrycks, D., Schmidt, E., and Wang, A. 2025. Superintelligence Strategy.

Herman, A., 2013. Freedom's forge: How American business produced victory in World War II. Random House Trade Paperbacks.