Expert Insights
PERSPECTIVE ON A TIMELY POLICY ISSUE

MICHAEL J. MAZARR, ANCA AGACHI

Building a U.S. National Strategy for the Artificial Intelligence Era

Drawing Lessons from Six Archetypal Strategies

November 2025

For more information on this publication, visit **www.rand.org/t/PEA4297-1**.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.
© 2025 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, visit www.rand.org/about/publishing/permissions.

PE-A4297-1

About This Paper

In this paper, we seek to inform the possible contours of a U.S. national strategy for the artificial intelligence (AI) era. We survey existing global AI strategies, examine nine major challenges to AI strategy development, and focus on assessing six archetypal strategies for insights and lessons. We conclude by outlining essential major components of any U.S. AI strategy and suggest an initial set of nine practical policy actions that can serve as foundation stones. Our aim is to support U.S. policymakers and members of the research community as debates about U.S. AI strategy continue.

Technology and Security Policy Center

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This work was undertaken by the division's Technology and Security Policy Center, which explores how high-consequence, dual-use technologies change the global competition and threat environment, then develops policy and technology options to advance the security of the United States, its allies and partners, and the world. For more information, contact tasp@rand.org.

Funding

The Geopolitics of AGI Initiative is independently initiated and conducted within the Technology and Security Policy Center using income from operations and gifts from philanthropic supporters, which have been made or recommended by DALHAP Investments Ltd., Effektiv Spenden, Ergo Impact, Founders Pledge, Charlottes och Fredriks Stiftelse, Good Ventures, Jaan Tallinn, Longview, Open Philanthropy, and Waking Up Foundation. A complete list of donors and funders is available at www.rand.org/TASP. RAND donors and grantors have no influence over research findings or recommendations.

Acknowledgments

The authors are grateful for over a year of detailed discussions in the Geopolitics of AGI Initiative and in the wider Technology and Security Policy Center at RAND, which informed this research, as well as for the dozens of colleagues who shared their wisdom on relevant topics. We are especially grateful for the support of Jim Mitre and Joel Predd in this process.

We are thankful to the RAND workshop participants, especially those who provided detailed feedback on this document: Joel Predd, Jim Mitre, Colin Kahl, Salman Ahmed, Matan Chorev, Beba Cibralic, Iskander Rehman, Helen Toner, Ben Buchanan, Gregory Smith, and others. We are grateful to Salman Ahmed and Ben Buchanan for thoughtful reviews that helped us to improve the final product. We thank David Glickstein for his contributions to the infographics.

Contents

About This Paper iii
Figures vii

Building a U.S. National Strategy for the AI Era: Drawing Lessons from Six Archetypal Strategies 1
National AI Strategies: A Crowded Field 2
Developing Strategies: Nine Key Challenges 4
Deriving Insights from Six Archetypal Strategies 12
Crosscutting Themes and Insights Across the Strategies 20
A Comprehensive U.S. AI Strategy: Discovering the Foundation Stones 22
Conclusion 28

Abbreviations 29
References 31
About the Authors 33

Figures

1. Ranges of Uncertainties in AI Model Development 8
2. Strategy 1: Preserve Responsible U.S. AI Technology Leadership 14
3. Strategy 2: Accelerate Domestic Diffusion and Application 15
4. Strategy 3: Establish U.S. Global Network Predominance 16
5. Strategy 4: Ensure Domestic Safety and Alignment 17
6. Strategy 5: Maximize Shared Benefits for the American People 18
7. Strategy 6: Maximize Shared Benefits for the World Community 19

Building a U.S. National Strategy for the AI Era: Drawing Lessons from Six Archetypal Strategies

The Industrial Revolution was the most geopolitically disruptive economic and technological transition in history. It prompted astonishing, unprecedented leaps in economic growth, living standards, and industrial capacity. Once the revolution hit, some countries had the inherent social, economic, and political characteristics that allowed them to prosper—with the United Kingdom (UK) at the top.[1] One—Japan, in the Meiji era—recognized the transition in progress and staged a dramatic catch-up effort. In the end, the countries that mastered this revolution bulled their way to the leading places in world politics, whereas those that could not, from the Ottoman and Austro-Hungarian empires, to China, to Russia, were left behind.

Imagine if the national leaders of Germany, Japan, the United States, and dozens of other countries had been given advance warning of the Industrial Revolution, as well as the chance to decide their fate. What would they have done? What strategic choices would they have made to flourish in the new era?

That is exactly the situation that the world finds itself in today. As of this writing in late 2025, the United States is on the cusp of another transformative techno-industrial transition: the Artificial Intelligence (AI) Revolution. But this time, there *is* advance notice. Although the ultimate shape of the transformation remains uncertain, there is little question that the AI Revolution will become a potent new sorting mechanism in geopolitics, separating a new set of leaders from the laggards and requiring advanced strategic thinking to make wise choices for their countries' futures. This effect may take more time to emerge than some AI enthusiasts claim—a decade or more rather than the very few years proposed by some. However, the full application of AI to thousands of business, government, education, military, and other tasks offers the potential for significant increases in productivity, growth, efficiency, and, ultimately, national power.

The dawn of the AI age will likely reshape power relations among countries. As AI's leading models approach levels of artificial superintelligence (ASI), AI has the potential to turbocharge national productivity and growth, scientific progress, and military capacity. At the horizon of our imagination, ASI may even reorganize the building blocks of human societies and functioning. As in prior transitions of this scale, nations that master the transition to a new paradigm—not only the development of AI models but also the supporting infrastructure, societal reform, and change that enable it—may well be the leading powers of that era. To emerge triumphant, countries require a strategic approach.

It is thus hardly surprising that more than 75 countries have developed national AI strategies.[2] In July 2025, the Trump administration issued its long-awaited AI Action Plan, which includes a wide-ranging outline of short- and medium-term steps that are needed to preserve U.S. AI advantages.[3] Yet even with that release, relatively little attention is being paid to the challenge of developing a truly comprehensive U.S. national strategy for success in the AI era. Most plans and proposals, including the AI Action Plan, focus on staying ahead in the technology itself rather than on any broader concepts of societal transformation and competitiveness. Building strategies for the development and deployment of AI remains a critical task and an appropriate subject for policy agendas. But plans that focus on this goal do not fully consider the challenge posed by the technological revolution unfolding before us.

---

[1] Michael J. Mazarr, *The Societal Foundations of National Competitiveness*, RAND Corporation, RR-A499-1, 2022.

[2] Stanford University Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2024*, Stanford University, 2024.

[3] Executive Office of the President, *Winning the Race: America's AI Action Plan*, White House, July 2025.

To address this gap, we seek to advance understanding of the requirements of a national strategy for AI and its ultimate expressions, including artificial general intelligence (AGI) or ASI.[4] In this paper, we draw lessons gleaned from our survey of existing AI strategy documents, identify major challenges in developing a U.S. AI strategy, and offer a set of six archetypal strategies to help advance comprehensive U.S. strategic thinking. We define each of those strategies in detail according to a common template of factors and then derive common and integrative themes from that list. We conclude by identifying several no-regret steps that can serve as the initial foundations of a comprehensive strategy.

Our suggestions should be seen as a complement rather than a corrective to the Trump administration's AI Action Plan. That plan represents a truly important step along the road to a broad-based U.S. AI strategy, in part through its emphasis on several themes, including the importance of open AI models and government support for AI diffusion and application. However, as its name indicates, the Action Plan is rightly focused on the most-important short-term steps to sustain U.S. technological leadership in the development and use of AI. Many other supporting actions will be needed to translate technological primacy into a durable national advantage; as the Action Plan itself states, that advantage must be achieved in ways that benefit American workers, the American people, and humanity—which is also the stated mission of the leading AI developers. The combination of these objectives, along with others, would be at the heart of a comprehensive AI strategy, which we also aim to develop through this paper.

Our thinking was informed by four sources of evidence and ideas. First, for the past 18 months, we have been engaged deeply in RAND's Geopolitics of AGI initiative by reading hundreds of essays; hosting dozens of meetings with leading thinkers from academia, government, and firms developing leading AI models; and conducting fresh research within RAND. Second, we reviewed the AI strategies that are emerging around the world, using the comparative assessment as a baseline for this perspective. Third, we conducted a workshop with experts in AI and geopolitics, which was held at RAND in June 2025. Finally, both as an experiment and a separate line of research, we also employed large language models (LLMs; specifically, Claude Opus 4, Sonnet 4, and ChatGPT 4o, which were the leading public models as of this writing in June 2025) to generate ideas across all aspects of this research through the definition of strategies, identification of crosscutting themes and findings, and suggestion of policy responses. We integrated many of these ideas into our thinking.

National AI Strategies: A Crowded Field

Recognizing the growing importance of AI, dozens of countries have laid out versions of what might be termed a national AI strategy.[5] As of 2024, 75 countries have released such strategies, with Canada being the first to do so in 2017.[6] Major economies and technology-forward countries led the initial wave of published strategies (circa 2017–2021), whereas countries in the Middle East, Africa, and the Caribbean created their documents in later years (circa 2021–2025), showcasing the breadth and depth of AI expansion. Beyond

[4] There are numerous debates in the field concerning the definition of AGI. For the purposes of this paper, *AGI* is defined as an AI system that is "at least as capable as a human at most tasks." For more on definitions of AGI, see Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg, "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv, arXiv:2311.02462v4, June 5, 2024, p. 1.

[5] This analysis considered specifically documents labeled as "strategies." Although these documents are representative, a more expansive analysis should consider additional documents (such as plans, memorandums, and funding levels) to provide a comprehensive picture of national goals and approaches to navigating the AI era. During the time in which this report was developed, new frameworks were released that wrestle with the rapidly evolving nature of the technology and its increasingly transformative implications.

[6] Government of Canada, "Pan-Canadian Artificial Intelligence Strategy," webpage, last updated December 3, 2024.

national AI strategies, there are numerous associated national documents (e.g., principles, guidelines, plans, memorandums, white papers) that provide the policy and legislative foundation for many countries around the world. Strategies directed specifically toward superintelligent AI or AGI are still rare, even as more documents ponder the implications of a rapidly advancing general-purpose technology.[7]

The existing national AI strategies are also presented in highly heterogeneous documents that vary widely in terms of focus, directionality, and associated resourcing—ranging from sector-specific approaches to whole-of-government (and, in some cases, whole-of-society) efforts. Many documents act as guiding principles and provide a general vision for a country's (or supranational institution's) place in an AI-shaped world, whereas others offer clear strategic directionality and a "comprehensive approach merging socio-political and economic incentives."[8]

The strategies display crosscutting themes. In the following list, we highlight several common themes, as well as examples:

- *AI innovation and research and development (R&D).* AI innovation is considered an important part of several strategies (Canada, the Republic of Korea [ROK], China) and is seen as being directly connected to and fostered by R&D investments within an integrated, mutually reinforcing ecosystem.
- *AI adoption and commercialization*. AI adoption and commercialization are considered important to ensure that economies maximize productivity and gains, as well as maintain pace with the ongoing transformation (European Union [EU], China). There are varying levels of specificity concerning the exact sectors targeted, and a few strategies also consider the potential impact or upheaval that would be elicited by large-scale AI adoption. The integration or mainstreaming of AI—especially frontier AI—across the public sector and national security enterprise is also featured in some strategic documents (United States). More-recent documents also increasingly discuss the enabling infrastructure that is necessary for AI's adoption or commercialization, such as compute (ROK, EU), energy (United States), data, and talent (see the following bullet).
- *AI talent*. All strategies discuss the importance of attracting, maintaining, training, and growing top-tier AI talent. Although some strategies focus on attracting AI talent in research, academia, and industry (Canada, United States), others consider AI (up)skilling to be an important part of the broader commercialization of AI to ensure that no one is left behind (EU).
- *AI governance*. AI governance, including ethics, safety, and trustworthiness, feature heavily across multiple strategies.
- *AI cooperation*. Most strategies recognize the critical role of industry in securing AI leadership, and some even discuss the importance of public-private partnerships in broader industrial development.

Some of the national AI strategies surveyed display specific focus areas, revealing different conceptions of how AI may transform society. For example, both the EU and Japan explicitly center humans in their AI national strategies, and Japan highlights the importance of human dignity and sustainable development. Most strategies discuss the importance of adopting and mainstreaming AI for economic purposes; however, other strategies, such as those from China and the United States, also discuss the military applications of AI, but only the latter mentions the importance of appropriate governance of these efforts.

[7] Dan Hendrycks, Eric Schmidt, and Alexandr Wang, "Superintelligence Strategy: Expert Version," arXiv, arXiv:2503.05628v2, April 14, 2025.

[8] Roxana Radu, "Steering the Governance of Artificial Intelligence: National Strategies in Perspective," *Policy and Society*, Vol. 40, No. 2, June 2021, p. 184.

Although A(G)I is still a rare focus in most national strategies, there is growing awareness about the general-purpose nature of AI and the risks that it poses. In addition to acknowledging the general-purpose nature of AI, more-recently released documents (since 2021) acknowledge its potential for leading to AGI and the associated risks that this evolution may pose. As early as 2021, the UK's AI strategy urged readers to "take seriously the possibility of AGI" and consider long-term risks.[9] More recently, such documents as the EU's *AI Continent Action Plan* and the White House's memorandum "On Advancing United States' Leadership in Artificial Intelligence" recognize the increasing general nature of frontier AI models and the potential associated risks and opportunities that this technology poses.[10]

The strategies surveyed note the need for leadership in AI: The UK wishes to "remain an AI and science superpower fit for the next decade";[11] the United States aims to maintain leadership globally in "the responsible application of AI to appropriate national security functions";[12] the EU wants to "be the champion of an approach to AI that benefits people and society as a whole";[13] and of course, China hopes by 2030 to become "the world's primary AI innovation center" and a "leading innovation-style nation and an economic power."[14] This leadership role is not unique to major powers: Such countries as Australia, Canada, and ROK also seek to carve out similar leadership roles for themselves.

Our preliminary assessment suggests that most governments and research communities understand the transformative nature of AI (even while the thinking concerning AGI is still nascent) and that some common elements of importance in securing national AI-enabled adaptation have been identified. However, the strategies are not sufficiently specific regarding the *prioritization, trade-offs, risks, and societal relationships* that are necessary for success, nor do they detail the socioeconomic incentives that are necessary to manage the transition (for instance, in the case of large labor displacement). More broadly, the strategies focus more on the *technology itself* and the input required for success (e.g., talent, infrastructure) and less on the *societal impact that A(G)I will have* and the *outcomes* the strategies will create. In a nutshell, if the transformation ushered in by AI will be as tectonic as expected, a provocative thought is that a national AGI strategy might need to provide more than just discrete actions: Rather, it may need to provide a roadmap for the large-scale transformation of human societies.

Developing Strategies: Nine Key Challenges

All these initial national strategies reflect tentative assessments at the start of an era that is fluid and rapidly emerging. They also largely embody broad statements of intent rather than detailed assessments that grapple with the complexities of the coming transformation. As the United States and other nations continue to

[9] Government of the United Kingdom, *National AI Strategy*, version 1.2, September 2021.

[10] European Commission, "The AI Continent Action Plan," webpage, April 9, 2025; Joseph R. Biden, "On Advancing the United States' Leadership in Artificial Intelligence; Harnessing Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence," memorandum, White House, October 24, 2024.

[11] Government of the United Kingdom, "The UK's National AI Strategy," flow diagram, undated.

[12] Biden, 2024.

[13] European Commission, "Communication from the Commission: Artificial Intelligence for Europe," COM(2018) 237 final, April 25, 2018.

[14] Graham Webster, Rogier Creemers, Elsa Kania, and Paul Triolo, "Full Translation: China's 'New Generation Artificial Intelligence Development Plan' (2017)", DigiChina, Stanford University, August 1, 2017.

develop their strategies, bringing them into greater alignment with their national priorities and goals, they will have to deal with the implications of several major challenges.

Challenge One: Defining Clear yet Enduring Objectives

Architects of any strategy need to begin by defining its key objectives. With AGI, the difficulty is deciding the scope of the effort. Is the United States seeking a national competitive position in AI technologies or a wider national—and even global—pursuit of success in the AI era? So far, most U.S. strategy discussions, policies,[15] and proposals have focused on building up the technological foundations to achieve a competitive position,[16] enabling success without defining it.

But a strategy aimed at technological superiority will not be enough for at least three reasons. First, succeeding in a techno-industrial revolution requires that attention be paid to broader economic and social factors, such as workforce quality, labor conditions, supporting infrastructure, and regulatory systems. Second, any strategy has to embrace governance as a means of managing safety and alignment issues: AI that is out of control will not produce advantage. And third, any strategy must position the United States as the leader in making the AI era work for humanity by sharing its global benefits; if not, the United States is likely to face blowback and instability from places left behind.

We suggest six baseline objectives for U.S. AI strategy, which represent a mix of narrow technological goals and broader aspirations for national and global success in the AI era. The United States should

1. sustain a responsible leadership role in the development of AI capabilities, as well as the technologies and capacities directly associated with them
2. establish U.S. models, applications, and standards at the core of emerging global AI networks and ecosystems
3. lead a global effort to ensure AI safety and alignment, establish response strategies for loss-of-control incidents, and mitigate the impacts of the ongoing technological transition
4. catalyze the application of AI throughout the U.S. economy, society, and institutions
5. maximize the shared benefits of the AI transition for the American people in ways that reinforce individual and collective agency and dignity, as well as maintain the democratic character of the nation
6. support the responsible development and application of AI globally, including through the assistance of other countries in diffusing its use and maximizing its shared benefits to address some of humanity's most-pressing concerns.

Later in this paper, we offer an archetypal strategy that prioritizes each of those six goals and then stress test each strategy for vulnerabilities and gaps. The message here is simple: Succeeding in any new socioeconomic era demands far more than technological superiority,[17] and any national AI strategy must take this fact into account.

Discussions at the RAND-led June 2025 workshop highlighted an important aspect of this challenge: The United States should be specific about what it is trying to achieve. Many terms that are commonly used in such objectives—such as dignity, agency, equality, dominance, or advantage—are highly contested. The idea of ensuring that AI works to the benefit of all Americans is easily stated but has multiple meanings, many of

[15] Executive Order 14110, "Advancing United States Leadership in Artificial Intelligence Infrastructure," Executive Office of the President, January 14, 2025.

[16] Department for Science, Innovation & Technology, *AI Opportunities Action Plan: Government Response*, Government of the United Kingdom, 2025.

[17] Mazarr, 2022.

which have embedded, normative assumptions. Many seemingly obvious objectives are filled with tensions and dilemmas. When we speak of empowering human dignity, for example, do we mean for everyone regardless of their effort, or is it enough to guarantee dignity for those who work to make themselves successful in this new era?

Challenge Two: Anticipating the Trajectory of AI Development

Beyond defining goals, the architects of a U.S. AI strategy must confront a second challenge related to the character of the technology itself and the patterns of its emergence over the coming decade. The speed and character of AI-model development could take divergent paths based on such factors as the technology's generalizability and level of automation.[18] The United States might need very different strategies depending on those outcomes.

Take AI-model capability: Expert opinions differ on how reliable and powerful AI models are and will be in the coming years. This debate was renewed with the release of ChatGPT 5 in August 2025, which led some unimpressed observers to proclaim that generative AI scaling had hit a plateau and to suggest that AGI remains many years off. A strategy for a more gradual arrival of AI models with many limits might require very different steps from one that lays the groundwork for ASI. Arvind Narayanan and Sayash Kapoor argued in April 2025 that AI is a "normal technology" that will behave like previous technology revolutions,[19] playing out in decades, whereas many AI lab leaders continue to claim that its transformative applications are around the corner. A national strategy for a decades-long transition would not look the same as one that is designed for a quicker revolution.

A similar question surrounds another aspect of AI development: the issue of a potential takeoff point. Some experts believe that, once AI models begin to substitute for AI researchers (at huge numbers—10,000 or 1 million such experts), AI progress will hit an inflection point and rocket upward.[20] Any countries that are not on that rocket will be left far behind. Not everyone agrees that a takeoff is near or even plausible, but a strategy that anticipates such a takeoff point would lead to different choices in contrast to one that downplays such a possibility.

Furthermore, many researchers and leaders of AI labs believe that AI development is set to cross dramatic thresholds by 2027 at the latest,[21] thresholds that will deliver models that match or exceed the most-advanced human cognitive capabilities in many fields, including AI research. Although other researchers doubt whether such a takeoff is imminent, any U.S. AI strategy must be robust against the possibility of the rapid emergence of any version of AGI.

A third technology issue concerns open versus closed or proprietary AI models. The future AI landscape could be dominated by the use of open-source or open-weight models (such as Llama, Mistral, or DeepSeek) that can be copied, used, and, in some cases, modified by anyone. Such a world would call for very different approaches from one in which proprietary, AI lab–controlled models (such as the various incarnations of

[18] William Marcellino, Lav Varshney, Anton Shenk, Nicolas M. Robles, and Benjamin Boudreaux, *Charting Multiple Courses to Artificial General Intelligence*, RAND Corporation, PE-A3691-1, April 2025.

[19] Arvind Narayanan and Sayash Kapoor, "AI as Normal Technology," *AI as Normal Technology* blog, April 15, 2025.

[20] Leopold Aschenbrenner, "Situational Awareness: The Decade Ahead," *Situational Awareness* blog, June 2024.

[21] The CEOs of Open AI (Sam Altman), Google DeepMind (Demis Hassabis), and Anthropic (Dario Amodei) all estimated that AI may surpass human capabilities by 2027. For more, see Josh Tyrangiel, "Sam Altman on ChatGPT's First Two Years, Elon Musk, and AI Under Trump," Bloomberg, January 5, 2025; Alex Kantrowitz, "Google DeepMind CEO Demis Hassabis: The Path to AGI, LLM Creativity, and Google Smart Glasses," *Big Technology* blog, January 23, 2025; Berber Jin and Joanna Stern, "Anthropic CEO Says AI Could Surpass Human Intelligence by 2027," *Wall Street Journal*, January 21, 2025. See also Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, "AI 2027," AI Futures Project, April 3, 2025.

ChatGPT, Claude, or Gemini) provide the lion's share of applications. To make matters more challenging, the issue of open models is immensely complex: The debate about open versus proprietary models may sound very different depending on the application, the user, the time frame, and many other variables.[22]

The future of open models overlaps with a fourth and equally important issue about AI models: what so-called commodified AI can achieve as opposed to frontier models. There are advocates who believe that the big models will be better at everything than less-advanced ones and that the firms developing them will dominate the vast majority of AI applications. In this scenario, firms that want to use AI for logistical planning, hospitals that hope to use it for diagnostic or administrative purposes, and other similar "customers" will all eventually turn to the Anthropics and OpenAIs of the world because their models will do a better job.

Other researchers have looked at the collapsing costs of AI models that were just recently at the frontier, which are now six months or a year behind the most-powerful models, and have taken into account the rise of open models and a proliferation of smaller, cheaper, bespoke models developed for specific sectors to forecast a kaleidoscopic future in which hundreds of different models are in use. A midsize firm that needs AI for some task might be served nearly as well by a cheaper model built on the firm's own training data as by a frontier model. When AI becomes a commodity, the competitive and strategic future looks different from one dominated by a handful of models. Many other issues that affect AI technology will shape AI strategy, including the kinds of training and development methods used to shape models, the degree to which training is centralized versus distributed, how quickly *agentic* AI emerges (which is capable of operating independently of human guidance and oversight), the global spread of AI capabilities, and the ability to deliver on the growing energy needs of AI.

Any AI strategy needs to be robust against many possible combinations of these outcomes. Several of these issues have already confronted U.S. policymakers with near-term decisions regarding the degree to which the spread of open models should be promoted (or allowed for), whether to encourage or prevent the wider spread of capable AI models in the hands of many actors, and whether to race toward ASI.

An important related issue is the question of *model distillation*—the potential to get smaller models to mimic the outputs and capabilities of larger ones. Distilled models can be 80 to 90 percent as capable as a bigger version while being several times smaller and faster. Researchers have already had some success with this approach. To the extent that model distillation succeeds more generally, it would put more-capable models in the hands of more people and into more edge applications. If U.S. firms remain at the forefront and shape the process of distillation from their models, competitive advantage might be sustained. But there are many complexities in this process, and, as distilled models become more common, the strategic implications could run in many directions.

Figure 1 depicts a few of these ranges of outcomes and describes a few of the many combinations that could emerge. U.S. AI strategy needs to be robust to many of these mixtures and emergent patterns. Most importantly, as we have noted, it needs to be resilient to the most important potential development: an explosion of the capability of AI models that represents a breakthrough approaching AGI. As we will argue, apart from more serious concerns about AI alignment and control, many elements of a baseline strategy would likely remain the same even for an "AGI breakout" scenario. But the urgency of a coherent strategy would become vastly greater.

[22] Jon Bateman, Dan Baer, Stephanie A. Bell, Glenn O. Brown, Mariano-Florentino (Tino) Cuéllar, Deep Ganguli, Peter Henderson, Brodi Kotila, Larry Lessig, Nicklas Berild Lundblad, Janet Napolitano, Deborah Raji, Elizabeth Seger, Matt Sheehan, Aviya Skowron, Irene Solaiman, Helen Toner, and Polina Zvyagina, *Beyond Open vs. Closed: Emerging Consensus and Key Questions for Foundation AI Model Governance*, Carnegie Endowment for International Peace, July 2024.

FIGURE 1
Ranges of Uncertainties in AI Model Development

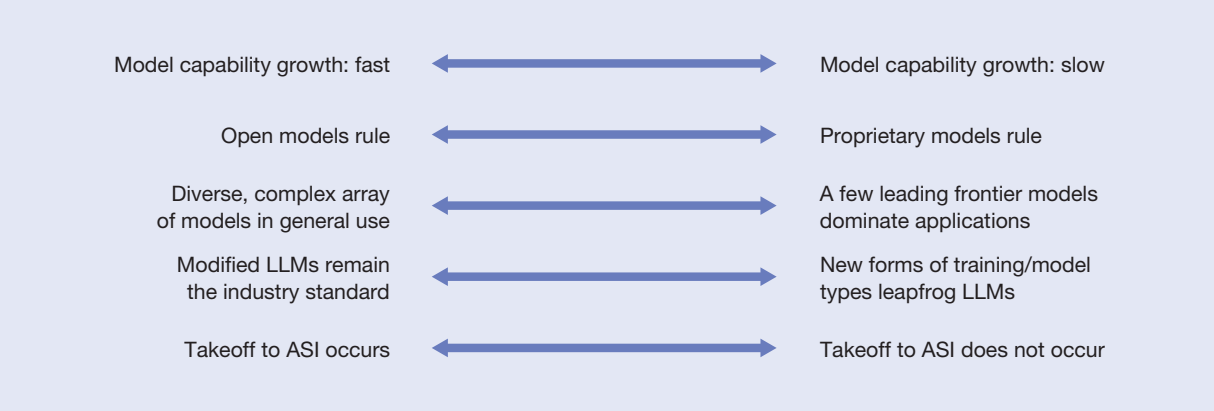

NOTE: These uncertainties could play out in complex mixtures with different strategic effects. Examples include

- a combination of rapid model capability growth within modified LLM architectures with the rise of open source to become more commonly used
- takeoff to ASI based on just one proprietary model that dominates by representing a breakthrough in new model training and design approaches
- a radically decentralized world of dozens of near-frontier models, open and closed, mixing LLMs and other approaches.

Challenge Three: Defining Success as More Than the Technology

One lesson of history is that the countries that race ahead in techno-industrial revolutions are not always those whose industries and investors come up with the most astounding discoveries and new technologies. The nations that eventually dominate such transformations do so because they possess societal qualities that spur the productive development, deployment, and use of the new technologies and the ability to adapt socially and economically to the demands of the new era. A strategy for the Industrial Revolution circa 1800 that focused narrowly on building the best steam engine would not have delivered lasting competitive advantage: Wider social, institutional, resource, and financial issues made the critical difference for leading nations.

Part of the lesson here is that strategies for *diffusion* (i.e., deploying and using a technology throughout an economy and society) are as important as—if not more important than—developing the technology itself,[23] as the UK's success compared with that of France during the Industrial Revolution shows. In other words, the United States needs to develop a whole-of-society strategy. AI must be thought of as one *means* to produce a dynamic, highly competitive society—not the single source of economic success. A comprehensive strategy for success in the AI era will have to include such things as support for broader innovation and R&D in AI applications; programs to develop and attract talent; campaigns to reform public- and private-sector organizational structures to employ AI capabilities that deliver efficiencies, as well as dignity and autonomy, to people who work in and interact with them; and much more.

Flourishing in the AI era is not merely about building the best model, any more than success in the Industrial Revolution was merely about having the best spinning jenny or coke-smelting process. It is about having the most dynamic, adaptive, and resilient *society*. As ambitious as it may be, any meaningful strategy for success in the AI era must address these larger issues. This perspective will become even more impor-

[23] Jeffrey Ding, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, Princeton University Press, 2024. See also Jeffrey Ding, *Running the Right Artificial Intelligence Race: A National Strategy for AI Diffusion*, RAND Corporation, PE-A4165-1, September 2025; Radha Iyengar Plumb and Michael C. Horowitz, "What America Gets Wrong About the AI Race," *Foreign Affairs*, April 18, 2025.

tant the closer AI models get to ASI, when their socioeconomic effects are likely to be powerful and potentially destructive.

Challenge Four: Dealing with the Public-Private Dilemmas

Any national AI strategy must also confront the immense challenge of managing the public and collective effects of a technological revolution whose conception, development, and deployment are dominated by the private sector. As of this writing, the majority of decisions are being made by private firms—including AI labs and other private-sector companies adopting and tweaking AI models—as well as individual citizens (such as people who take open models and retrain or reprogram them for specific purposes). In many cases, government efforts to limit, control, or shape those outcomes will be difficult and controversial. This distinction and gap between the public and private sectors is one of the defining elements of the AI strategy challenge and likely requires an approach that is distinct from classic national security strategies that have long been in the province of governments.

In an era of national competition in AI, the public-private relationship will become increasingly important as AI developers ask for support to accelerate their progress and thus win the larger competition. For example, energy regulation, permitting, and construction exemplify areas in which this relationship will likely be both consequential and contested. There are places where government and some parts of the AI industry may end up on opposite sides of a specific debate. One challenge will be for government and business—especially the AI labs and leading firms in associated technologies—to develop a combined strategy for AI success. The goal should be to focus on shared interests and the potential for a powerful, mutually supportive relationship. Such collaboration will be difficult: AI's dual-use nature, ubiquity, and development almost exclusively in the hands of the private sector make the public-private relationship in the case of AI unprecedented.

It is not in the interest of AI actors in the U.S. government or in the private sector for there to be mass public backlash against the rapid development and deployment of AI because of unexpected spikes in unemployment, dangerous effects of AI bots (especially on young people), crippling cyber incidents, or the AI-enabled creation of novel pathogens in the hands of terrorists. These AI actors have a shared interest in anticipating and addressing potential public harm and consequent adverse public reactions that would undermine trust in government and hurt these firms' bottom lines. Yet the U.S. government and private sector are only at the earliest stages of taking such measures in a systematic fashion, and the U.S. government itself is hampered by a lack of expertise, dedicated capacities, and new processes designed for these purposes. Again, the absence of such resources indicates that government and business—both AI labs and major firms applying AI in their day-to-day work—must develop a deep and ongoing partnership to ensure that AI benefits the nation.

As a society, the American people need to decide what sort of future they want AI to create rather than merely watch as AI generates a future they may not like. But to shape and control AI in such a manner, the American people will need to act through mechanisms of collective interest. AI's potentially dramatic outcomes signal that governments will need to do much more to direct AI's growth and usage than they are doing today. The closer the world gets to AGI or ASI, the more this tension between seizing the value of AI and guarding against its risks will emerge.

Challenge Five: Considering the Potential for Radical Futures

A fifth challenge in developing AI strategy is the risk underplaying the mind-boggling outcomes—both positive and negative—that AI could bring about over the coming decades. Again, the central issue is AGI or ASI: By 2026 or 2027, according to some warnings from AI researchers, the capabilities of AI models could

become so powerful that their effects—in terms of destroying careers, manipulating information, and undermining social capital—would threaten the social and political foundations of countries.

Technical accounts of AI's potential now suggest that some of science fiction's most astonishing—and worrying—predictions might come true. AI could replace most human jobs. It could power armies of humanoid robots that are integrated throughout our society. AI models could fashion new ideologies and religions that rival human traditions. Millions of Americans might come to rely on AI-powered chatbots for companionship, friendship, and therapy. Many people might work with more AI agents—and perhaps count more of them as friends—than they do human beings. The AI era cannot be assumed to be an identifiable version of the present day, with slight changes; it could inaugurate a dramatically different form of human society, and any national AI strategy has to grapple with this possibility.

This challenge redoubles the importance of AI safety initiatives: These more-radical futures will serve humanity only if increasingly powerful AI is aligned with human needs. But this challenge also underscores the urgent need for fundamental discussions about what the United States wants the human future to be and how AI could serve human goals in a world in which peoples' autonomy is enriched—rather than marginalized—by AI. The need for proactive imagination must be balanced with the importance of staying grounded in the reality of what the technology can achieve.

Challenge Six: Operating in the Realm of the Possible for the U.S. Government

A challenge for the architects of any AI strategy is to develop a set of actions that can be feasibly implemented by the U.S. government. Government AI policy is beset by a huge bureaucracy, internal debates that take a year or more to resolve critical issues, a lack of sufficient expertise, and the absence of broad-based capacity to perform many functions—including the effective enforcement of export controls. In the meantime, fiscal pressures on U.S. government activities increasingly constrain the actions it can take, and it is unclear whether the U.S. government will have the resources to execute any elements of an AI strategy with very big price tags.

The designers of a national AI strategy face a real dilemma. In conceptualizing an AI strategy, they cannot be limited by what the U.S. government can do today. The stakes are too high to simply accept existing governmental shortcomings as a given. At the same time, idealistic strategies that are not based on actual government capabilities are wishful thinking. Having a strategy without the capacity to execute it is pointless.

Events will help to resolve this dilemma. If and when AI begins to clearly cause layoffs or is used for attacks on the American people, the public will demand action in ways that cause government officials to override their reluctance and take positive action. In such circumstances, the risk of developing the wrong solutions with hurried, politically charged actions in crises is very great. Therefore, one urgent priority is to prepare for those moments by thinking through the kinds of events that could occur, what the best responses would be, and what capabilities or authorities the government will need to execute those responses.

Part of our message on AI strategy is that ways to become more determined and more decisive in overcoming existing barriers to action and shaping the future that AI will create are needed. In the final section of this paper, we propose several initiatives that should form the baseline components of any long-term strategy; although they do not presuppose every decision in a strategy, they will likely help under any plausible scenario of AI's emergence.

Challenge Seven: Seeing AI as Part of a Wider Techno-Industrial Revolution

AI is likely to have dramatic effects on economic and military prospects and aspects of social stability. AI may become the hub around which other technological advances will be made by making them possible and playing a critical role in their managing once deployed. But AI is not the only technology undergoing a renais-

sance. Rapid advances in biotechnology, quantum physics, new forms of energy, advanced manufacturing, and other areas are taking place alongside AI. U.S. AI strategy needs to account not only for advances in AI but also for the creation of an AI ecosystem that supports the safe and effective development of a wide variety of emerging technologies.

One emerging technology that will be critical in setting the context for AI's effects is robotics. The real-world capabilities of AI will always be limited if they remain constrained to cognitive functions. Dozens of U.S. and Chinese firms are working to develop various versions of next-generation robots, including humanoid ones (although there are debates about how quickly humanoid robots will be able to perform high-complexity tasks). This analysis did not deal with the issue of robotics but with an element of U.S. AI strategy that will have to deal with the linkages between these technologies.

Even more broadly, the next decade will be characterized by dramatic advances in biotechnology, energy, and other fields. Although such advances are beyond the scope of this analysis, AI will be related integrally to all of them; consequently, nations will gain a competitive advantage not only from advances in AI but also from the breakthroughs that come from its relationships with these other areas. Many of our basic ideas for strategy (which we present later in this paper) address this connection, at least indirectly, by providing for AI leadership and effective implementation in ways that would support its role in partnership with these other technologies.

Challenge Eight: Navigating Evolving Geopolitical Relations

An AGI strategy will need to reckon with and provide a pathway for the United States to navigate shifting geopolitical waters because the "rules of the game" may change as a result of broader AI adoption and transformation. Military power, economic potential, and shared values will remain foundational elements of geopolitical relations for the foreseeable future. Yet as the adoption of AI advances in the United States and in other countries, the shifting foundations of national power in the AI era will reshape patterns of collaboration and competition in international relations as well. Access to data may matter as much as access to markets; chip assembly lines could be the new "arsenal" of tomorrow's democracies; shared values may be shaped as much by AGI talent on the move as by elites schooled in the same institutions of the 20th century; and energy access and capacity could be the new currency that powers the AGI economy. The intensity of the shift will likely depend on the final technological capabilities of AI, but it is plausible that—short of a takeoff scenario—U.S. diplomats and strategists will need to navigate a patchwork of technological change in which the old and the new coexist and shape the geopolitical realities of the day. If a radically more-powerful AI suddenly emerges, national competitive prospects and relations among countries will be subject to radical turbulence and potentially unstable transitional pressures.

A U.S. AI strategy should anticipate how trade relations, defense partnerships, and political cooperation may evolve around certain AGI-related elements, as well as where and how these elements link with other established factors that influence geopolitics. For example, the United States may choose to advance global partnerships to deepen its semiconductor supply chain, construct data centers, garner data access, and supply and benefit from new energy sources,[24] which would change the structure of its alliance architecture. Existing alliances may recede into the background or coexist alongside new ones. China will be pursuing its own goals in this transforming world—including in reaction to U.S. choices—and a U.S. AI strategy will need to anticipate where points of contention may be.

[24] Janet Egan, "The Way Forward: AI Diplomacy with Compute Leadership," in G*lobal Compute and National Security: Strengthening American AI Leadership Through Proactive Partnerships*, Center for a New American Security, July 2025.

Ultimately, however, a U.S. AI strategy will need to define clear goals for U.S. engagement in the world, including a policy toolbox that is attuned to the demands of the time, as well as clear red lines.

Challenge Nine: Making Normative Choices About the U.S. Society and Economy

Discussions at the RAND-led June 2025 workshop pointed to a broader challenge than simply defining the goals of an AI strategy: The United States needs to decide what future it wants AI to create rather than watch passively as a future is manufactured that it may not want. A strategy will need to be specific in its goals, yet many terms that are commonly used in such objectives—such as dignity, agency, equality, dominance, or advantage—are highly contested and have different meanings depending on the context. The idea of ensuring that AI works to the benefit of all Americans is easily stated but has multiple meanings, many of which are embedded normative assumptions. Many seemingly obvious objectives also turn out to be filled with tensions and dilemmas. Does empowering human dignity, for example, mean doing so for everyone regardless of their effort, or is it enough to guarantee dignity for those who work to make themselves successful in this new era?

Increasingly, the question that should be asked is not *what might AI do to any given issue but rather what do we want it to do*? Such a mindset will raise difficult policy issues about the proper role of government in balancing between imposing needed controls and safeguards and allowing AI progress to deliver its promised benefits. A strategy will be directive in addressing these factors, whether explicit choices are made or decisions are simply delayed. Therefore, early consideration is critical.

Yet the government is not the only actor that needs to be part of this effort. Nongovernmental organizations, universities, and private-sector firms all need to decide for themselves what sort of AI-influenced future they want. During the Industrial Revolution, many business leaders eventually decided that controlling the rough edges of the process through reform, including treating labor as a political and social equal and granting significant benefits to workers, was a necessary and, ultimately, a modest price to pay. However, awareness of the necessity of reform was always incomplete and took decades to emerge. A much more rapid emergent sense throughout society of the positive steps will be needed to shape the future the United States wants.

Deriving Insights from Six Archetypal Strategies

Given all these uncertainties and the accelerated timeline on which governments have had to confront the AI issue—even in the past two years—the United States is not in a position to develop a comprehensive, final national strategy for succeeding in the AI era. Too much remains in flux, including the evolution of frontier models, the state of semiconductor production, and the stance of third parties regarding the AI competition. The nature of the strategic challenge cannot be captured in a classic national security strategy–style document.

Furthermore, given the rapid pace of change that is emerging, any U.S. AI strategy is likely to be one that favors evolution and change over premeditated planning and requiring the constant monitoring of trends and adaptation to new developments. Concepts of emergent or adaptive strategy match the nature of this strategic problem better than rigid, deliberate approaches.[25] An adaptive approach identifies clear goals and a broad theory of success but recognizes that the vast majority of strategic judgments will have to be made in process throughout a flexible and fluid response to events as they unfold. Such strategies also assume a high

[25] Michael Mazarr, "AI and the Challenge of Adaptive Strategy," *Geopolitics of AGI*, Substack, July 2, 2025.

degree of bottom-up decisionmaking and experimentation that is designed to test a variety of approaches to new problems or opportunities as they emerge. It is a fluid, responsive, open-minded conception of strategy.

Both the nature and the pace of the AI issue mean that, as of this writing, no one is in a position to draft anything close to a comprehensive and reasonably final document that lays out the next decade of U.S. AI strategy. But the limits of our understanding of what is to come does not mean that we cannot say anything or make any decisions. The importance of carefully assessing the strategic environment, deciding on core objectives, developing an initial theory of success (which will have to be subject to continual reassessment), and choosing a handful of initial moves that support success under any future scenario are all the kinds of steps that are even more important in a context of emergent strategy. In such a context, the risk is that a strategic actor that is put off by uncertainty leaves strategic choice and developments to chance.

Even in a fluid context, the U.S. government—as well as private-sector and civil society organizations whose decisions and initiatives will be critical to any U.S. AI strategy—needs to develop an initial strategic vision and set of actions from which it can then adapt. A critical part of that initial vision is identifying baseline or foundational policies or actions that will be robust in any future scenario and addressing the most-critical requirements for success. To inform that thinking, we outlined six broad, archetypal AI strategies, examined them in detail, and derived crosscutting insights by comparing and contrasting them.

For each of the strategies, we defined several component elements, including the basic goal of the strategy, relevant aspects of the global context, key assumptions, focal problems and opportunities of the strategy, more-detailed objectives, the strategy's theory of success, sources of U.S. competitive advantage in the strategy, specific policy actions, and trade-offs. In Figures 2–7, we draw on that thinking to lay out several key elements of each strategy: its basic concept, essential goal, theory of success, and selected views about key global trends, assumptions, risks, and challenges.

We decided to use the six basic objectives for U.S. AI strategy mentioned previously as the entry point for developing archetypal strategies. We could have used other approaches; there is an almost endless list of potential archetypal strategies for AI. (For example, one could use a scale of unilateral versus multilateral approaches to the challenge: Should the United States seek self-sufficient predominance to achieve its AI goals or link itself to a multilateral alliance of AI partners?) To some degree, the organization of categories is arbitrary. But it is important to recognize that this is only one way of identifying divergent strategic options.

The resulting strategies are not mutually exclusive. No set of archetypal strategies will be: Any overarching strategy will need to incorporate components that address all these objectives.[26] However, each of these strategies represents a specific emphasis and set of priorities, and debates about both emphases and priorities could help clarify larger strategic discussions. Each strategy offers a theory of success that contains an argument—either implicit or explicit—as to why that specific emphasis should be central to U.S. strategy and why its particular focus offers the catalytic effect essential to overall success. Each embodies ways of assessing and dealing with the challenges outlined earlier in this paper.

We used the detailed comparisons of these strategies to build two sets of findings. First, we highlighted key trade-offs and choices that any U.S. AI strategy must confront. Second, we identified several crosscutting themes and insights that can help guide U.S. AI strategy. The next section lays out findings in both those areas. With those findings in mind, we then turn to the task of building a framework for U.S. AI strategy.

[26] For a similar approach based on a different variety of AI futures, see Bengüsu Özcan, Daan Juijn, Jakob Graabak, and Sam Bogerd, "Advanced AI: Possible Futures—Five Scenarios for How the AI-transition Could Unfold," Centre for Future Generations, July 10, 2025.

FIGURE 2

Strategy 1: Preserve Responsible U.S. AI Technology Leadership

Strategy #1

Preserve Responsible U.S. AI Technology Leadership

This strategy emphasizes large private and infrastructure investments to reinforce U.S. and allied advantages in the AI stack, countering China's advances and ensuring that safety remains central to technological leadership.

What leading **global trends** inspire this strategy?

U.S. dominance

U.S. models and private-sector capital flows have a significant lead.

China seeks leadership

China seeks relative advantage through various means, especially diffusion strategies.

Third-country hedging

Third countries aim to partner with those who can help them harness AI benefits.

What is the strategy's **theory of success**?

This strategy rests on the logic **that dominance in the core technology stack creates cascading advantages**. Technological edge translates into economic advantages through higher-value industries and military advantages, ultimately preserving overall U.S. strategic dominance.

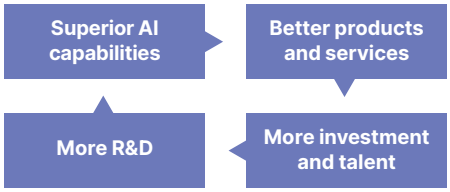

What are the **underlying assumptions** supporting the strategy?

Private-sector capital infusions continue; scaling techniques do not bottleneck; China or others do not discover radical new methods to improve model capabilities.

What are the **lines of implementation**?

Private-sector activities

Private capital and AI labs provide essential energy.

Government incentives

U.S. government supports private sector by removing regulatory barriers and providing economic incentives.

Public-private collaboration

Public-private partnerships achieve necessary infrastructure progress and sustain AI alignment.

International cooperation

United States cooperates with third countries within bounds to provide more capital and certain infrastructure elements.

Export controls

United States constrains cooperation with China and administers export controls to limit China's chip progress.

What **risks or challenges** could threaten the strategy's success?

Financial and resource restraints

U.S. government will be unable to replace depleted private-sector resources, partly because of budget restraints.

Technical and infrastructure limits

Roadblocks could impede model capability growth and infrastructure build-out.

Open model threat

Competitors could leverage open models to achieve success.

Chinese and global competition

China may quickly move to fund model application and use for developing nations, gaining access to capital and data.

FIGURE 3

Strategy 2: Accelerate Domestic Diffusion and Application

Strategy #2

Accelerate Domestic Diffusion and Application

This strategy prioritizes gaining strategic advantage through mastering and leading the diffusion, application, and use of AI to achieve economic, social, and military advantages rather than through model capability.

What leading **global trends** inspire this strategy?

U.S.-Chinese competition

While U.S. models still lead, China seeks advantage through diffusion strategies.

Third-country hedging

Third countries look for partners to help them advance in AI applications.

Need for effective governance

Amid global instability and emerging risks, governance effectiveness is a critical comparative advantage.

What is the strategy's **theory of success**?

This strategy rests on the logic that **mastering AI diffusion creates cascading competitive advantages**.

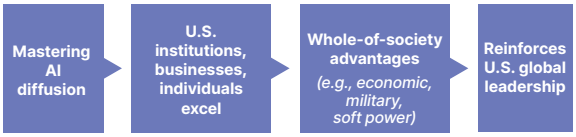

What are the **underlying assumptions** supporting the strategy?

Advantage comes more from AI diffusion than model capability; diffusion yields compound returns, such as first-mover gains, learning, standards, and influence; government can shape diffusion; power supply will meet inference demand; public trust will enable widespread adoption.

What are the **lines of implementation**?

AI integration

U.S. government integrates AI across all federal agencies and departments.

Government incentives

Federal, state, and local programs incentivize AI adoption to spread effective use cases throughout society.

Workforce retraining and education

Government establishes workforce retraining, education initiatives to prepare Americans for AI-augmented roles.

Favorable regulatory environment

Regulatory sandboxes and programs help accelerate private-sector AI experimentation.

Showcasing U.S. success abroad

International programs showcase U.S. AI successes to drive global partnerships and U.S. model adoption.

Safety regulations and oversight

Strong AI safety efforts ensure that AI remains aligned during diffusion process.

What **risks or challenges** could threaten the strategy's success?

Financial and resource restraints

Enormous costs and complexity of retrofitting existing systems and institutions for AI.

Limits on government control

Government cannot fully steer diffusion or broad adaptation given institutional, regulatory, and cultural barriers.

Adoption risks

Loss-of-control incidents, workforce disruptions, or significant AI failures will slow adoption or intensify resistance.

Unequal distribution of AI

Inequalities could worsen if AI benefits are not distributed broadly across society.

FIGURE 4

Strategy 3: Establish U.S. Global Network Predominance

Strategy #3

Establish U.S. Global Network Predominance

This strategy aims to establish the United States as the world's dominant hub for emerging AI networks, including AI models, data centers, and semiconductor chips.

What leading **global trends** inspire this strategy?

Global demand for AI

Global demand for AI services is surging, especially in application layers and user-facing interfaces.

U.S.-China competition

United States leads in cloud, application programming interfaces (APIs), and LLMs but faces pressure from Chinese platforms and open-source efforts, while China builds a parallel AI infrastructure.

Third-country hedging

Many developing nations exhibit deeply entrenched hedging behavior.

What is the strategy's **theory of success**?

This strategy rests on the logic that **U.S. AI platforms gain value through network effects and platform dominance**, creating a virtuous cycle that makes them indispensable to global AI adoption.

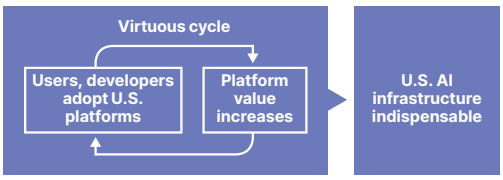

What are the **underlying assumptions** supporting the strategy?

Controlling central AI nodes shapes standards and flows; United States retains a capability lead; network centrality drives self-reinforcing dominance; users follow the most capable, most accessible platforms; first-mover advantages in AI infrastructure compound quickly.

What are the **lines of implementation**?

U.S. AI dominance

United States preserves its global lead in AI models and applications.

Financial incentives for global users

United States uses aid, diplomacy, and financial incentives to promote its AI platforms in emerging markets.

International cooperation

Strategic alliances with Europe, Japan, India, and Latin America help co-build U.S.-aligned AI stacks.

Trusted AI service networks

United States subsidizes/brands trusted AI service networks (e.g., in health, education, productivity) to drive adoption abroad.

Private-sector licensing

United States incentivizes private-sector licensing models that promote U.S. model use.

Favorable regulatory environment

Favorable regulatory environments help U.S. AI platforms expand internationally.

What **risks or challenges** could threaten the strategy's success?

Chinese competition

Chinese subsidies, bribes, and other means threaten to beat U.S. firms on price and influence.

Financial restraints

United States may lack the fiscal capacity for significant spending.

Backlash and resistance to U.S. AI

Rising global resistance to U.S. tech dominance may trigger regulations and digital sovereignty efforts.

U.S. control vs. global access

United States must control chip supply, although other countries will seek access to join AI networks.

FIGURE 5

Strategy 4: Ensure Domestic Safety and Alignment

Strategy #4

Ensure Domestic Safety and Alignment

This strategy centers on mitigating risks from misaligned AI, building response capacity for loss-of-control events, and positioning the United States as the global leader in safe AI, gaining reputational and practical advantages.

What leading **global trends** inspire this strategy?

Safety sidelined

Accelerating competition between states and AI labs has sidelined safety concerns; third parties develop models without attention to safety.

Weak global governance

Fragmented global governance for AI, with weak norms and limited enforcement capacity.

U.S.-Chinese competition

United States and China have mistrust and overall perception of zero-sum rivalry.

What is the strategy's **theory of success**?

This strategy imagines that **prioritizing safety and alignment while advancing AI creates sustainable competitive advantage**. Safe, trusted AI drives global adoption, network effects, and standard-setting. Responsible leadership attracts partners, builds coalitions, and reinforces U.S. influence.

What are the **underlying assumptions** supporting the strategy?

AI safety poses real security and welfare risks; responsible leadership attracts partners and advantages; safety-first development can coexist with technological edge; global cooperation on standards is achievable; loss-of-control events would hinder U.S. progress.

What are the **lines of implementation**?

AI safety institutes

Established & funded federal AI safety institutes have strong enforcement authority.

Safety regulations and standards

The U.S. government develops AI safety regulation and standards for private-sector companies.

Rapid response capabilities

The U.S. government develops rapid response capabilities for AI safety incidents or concerning developments.

Testing and certification frameworks

The U.S. government creates frameworks mandating model licensing, audits, and risk-tiered deployment thresholds.

International cooperation

The U.S. government develops AI safety partnerships and agreements.

What **risks or challenges** could threaten the strategy's success?

Technical challenges

The actual mechanisms for ensuring safety may be extremely difficult in practice.

Cooperation challenges

AI labs and international partners may be unwilling to cooperate; government may be unwilling to enforce standards.

Trade-offs with leadership and progress

Safety guarantees could slow U.S. AI progress and leadership; competitors achieve breakthroughs first.

Unequal legal enforcement

Ambiguity in alignment definitions could make implementation uneven or politicized.

FIGURE 6

Strategy 5: Maximize Shared Benefits for the American People

Strategy #5

Maximize Shared Benefits for the American People

This strategy seeks to accelerate the productive adoption of AI across society to enhance living standards, health, welfare, dignity, and autonomy while building a stronger, more competitive nation.

What leading **global trends** inspire this strategy?

U.S.-Chinese competition

The United States and China compete across social, economic, and military domains.

Public distrust in AI

Public increasingly distrusts AI systems, automation, and government competence in managing technological transitions.

Weak U.S. position

Governance effectiveness is a critical comparative advantage, yet U.S. domestic political context is fragmented.

What is the strategy's **theory of success**?

This strategy rests on the logic that **broad social buy-in and trust enable human-centered AI**, driving rapid, safe deployment and positioning the United States as a global model.

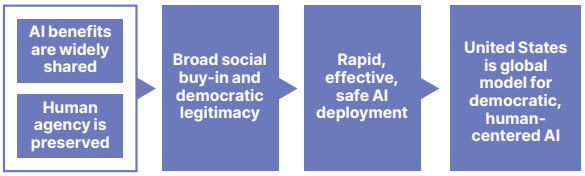

What are the **underlying assumptions** supporting the strategy?

AI's success requires clear benefits for Americans; legitimacy and social resilience underpin national strength; preserving human agency supports democracy; shared-prosperity strategies yield more lasting advantages than narrow tech or economic focus.

What are the **lines of implementation**?

Policies to ensure higher living standards

U.S. government establishes National AI Dividend Framework to ensure that economic gains from AI translate to higher living standards.

Community-centered AI pilots

Grants to civic groups, tribal communities, rural towns, and disadvantaged regions help localize innovation through tailored AI solutions.

Trust and protections

Enforce an AI Bill of Rights; ensure transparency, audits, explainability, public input, and relevant political reforms.

Workforce retraining and education

Investment in education and retraining helps Americans adapt to AI-augmented careers.

What **risks or challenges** could threaten the strategy's success?

Governance challenges

State and sectoral differences may slow coordination; federal mechanisms may lack agility.

Financial and technical restraints

High costs could affect comprehensive social support, redistribution efforts; technical restraints could challenge AI system development.

Backlash and resistance

Private firms may oppose equitable data use, worker voice, or profit-sharing.

Societal risks from AI

Rapid AI development could conflict with safeguarding social outcomes.

FIGURE 7

Strategy 6: Maximize Shared Benefits for the World Community

Strategy #6

Maximize Shared Benefits for the World Community

This strategy prioritizes supporting the responsible development and application of AI globally, such as by assisting other countries in diffusing and maximizing shared benefits to address humanity's most pressing concerns.

What leading **global trends** inspire this strategy?

Emerging global challenges

Climate change, pandemic preparedness, poverty, food insecurity, and sustainable development are growing in urgency.

AI competition and legitimacy

Unequal access and competition visions (authoritarian vs. democratic AI) influence international trust and alliances.

Third-country hedging

Third countries are interested in partnering with countries able to help them gain benefits from AI.

What is the strategy's **theory of success**?

This strategy rests on the logic that **leadership in leveraging AI to address global challenges generates many reinforcing advantages**.

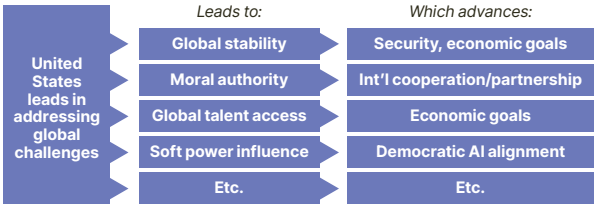

What are the **underlying assumptions** supporting the strategy?

Global prosperity and stability advance U.S. interests more than zero-sum competition; sharing AI benefits strengthens alliances and counters illiberal influence; humanitarian AI leadership builds lasting partnerships; U.S. AI leadership is trusted and welcomed.

What are the **lines of implementation**?

International AI research partnerships

United States establishes international partnerships focused on global challenges.

U.S. AI tools and platforms

U.S.-origin, open-source AI tools and platforms are designed for global development.

Assistance programs

U.S. funds development assistance programs to help nations build AI capabilities and global infrastructure projects.

International AI governance institutions

International institutions aimed at governing AI include broad global participation.

What **risks or challenges** could threaten the strategy's success?

Financial restraints

Global AI development assistance and infrastructure programs will likely incur enormous costs.

Chinese competition

China may underbid or subsidize partnerships with Chinese development and humanitarian initiatives, crowding out U.S. approaches.

Backlash and resistance

Governments or social actors in key countries and in the United States may oppose these initiatives.

Implementation challenges

Conflicts with national interests and diverse cultural, linguistic, and institutional contexts complicate deployment.

Crosscutting Themes and Insights Across the Strategies

A major purpose of this paper was to use a wide-ranging set of archetypal AI strategies to generate crosscutting insights and findings that can inform U.S. national AI strategy. In this section, we lay out several of our leading insights, which provide the basis for our suggested baseline elements of a national strategy.

Continued leadership in AI development is an essential precondition for all strategies. One lesson of the varying strategies was that some form of AI technology leadership plays a critical role in providing the basis for success in all of them. It is the essence of the first two strategies, but leadership is also necessary for networked power, which benefits the American people—or the globe—and ensures AI safety. The United States could still achieve some elements of several strategies without such leadership: It could share benefits among the American people, for example, regardless of its relative position in the race toward AI development. But the larger purposes of those strategies, such as setting a compelling example of AI's societal effects, can be fully achieved only if the United States remains at the forefront of AI development. Yet leadership can take different forms, including modeling capability; supporting the application and use of AI, safely; or undertaking policies to share benefits to humanity. The United States needs much better ideas of what degree of leadership is essential, effective, and possible. It could benefit from more-detailed analysis of what forms of leadership will be most relevant and how these different forms enhance or undermine each other. Ultimately, the United States may need to pursue different approaches based on needs, ranging from selective overmatching in some areas to cooperative engagement in others.

Becoming a world leader in domestic applications and use may be as important as leading in model capability. Strategies for *diffusion*—that is, demonstrating an ability to make AI work for the United States by promoting the effective application and use of AI models in practical settings throughout the economy and society—are essential to global leadership, efforts to share benefits domestically and globally, and the leveraging of intrinsic U.S. advantages, such as a free market and a decentralized, entrepreneurial social system. Domestic applications are the cutting edge of a practical AI strategy for enhancing national power; no strategy succeeds without their diffusion and use. One implication of this finding is that, at the moment, the United States has virtually no coherent national effort that is designed to promote AI diffusion and application.

Ensuring that the rollout and application of AI benefits Americans broadly and strengthens rather than weakens social solidarity are indispensable to competitive success. Such outcomes will help determine domestic stability, will, and governance, and the image conveyed by U.S. AI deployment will help shape international influence. A failure to demonstrate that AI is benefiting Americans, especially during a potentially massive AI-powered economic transition, would undercut to a significant degree any efforts to attract others to U.S. approaches. Any strategy needs to address the issue of promoting a better life, agency, and identity for the American people. Achieving such an outcome is just as important to true national competitiveness as leading in model capability or applications.

Major loss-of-control events—or significant malign uses of AI that cause harm—would create an existential risk to any strategy. Within one country or more broadly, an AI that demonstrates increasing loss of human control or is employed as a tool to harm people will trigger responses that cut against the competitive advantage being sought in any strategy other than the one that prioritizes AI safety. Safety is an essential enabler for all strategies. As a result, becoming a trusted leader of the AI development and deployment processes conveys significant power and influence. Demonstrating responsible, beneficial AI development creates sustainable competitive advantages by enhancing trust and increasing international cooperation.

Network advantages play a critical role in granting competitive advantage across all strategies, not merely the one that focuses on network power. Platform dominance and networks that are embedded globally will play a major role in creating self-reinforcing advantages. However, these advantages depend on existing U.S. networks of power, which are shaped by trade relations, political alliances, and cultural engagements. This

strategy is not as essential to national success as some of the others: If the United States leads in AI model development, diffusion and application, and benefit-sharing among the American people, it would gain a national competitive advantage and domestic solidarity. Over the medium- to long-term, however, a global AI ecosystem that becomes centered largely around Chinese models, hardware, and applications will create a dangerous context for U.S. power more broadly.

Many terms and concepts that are widely used in thinking about AI's effects are highly contested, stand in tension with one another, and do not provide guidance to AI strategy automatically. Such concepts as dignity, agency, or diffusion are important signposts, but their meaning changes in relation to their user and the context of use. Also, seemingly self-evident goals can be in tension. For example, maximizing the standard of living or pursuing economic growth may not overlap entirely with granting human agency and autonomy in the context of an economic system that has seen significant automation of human labor.

Governance effectiveness and quality that is powered, in some cases, by AI will be a critical differentiator in this new era. A combination of bureaucratic and institutional ossification, complex new problems, and political legitimacy issues contribute to a context in which the nations that can exercise effective, efficient, problem-solving domestic governance will gain a decisive strategic advantage. Increasingly, AI will empower such outcomes but may also undermine them. A campaign for escaping the trap of bureaucratic and political barriers to effective governance—newly empowered by AI—will provide important competitive advantages to any great power that can succeed in such an effort.

The relationship and competition between open and proprietary models could be a defining component of AI futures and needs to be a central focus of U.S. strategy. The contest between open and proprietary models emerged as a potentially critical determinant of success in many of the strategies—most notably in the strategy focused on global AI networks but also in those that are focused on diffusion, application, and bringing practical benefits to the American people and the world community. The United States needs to shape the emerging open model ecosystem and be a major player in it. This is an incredibly complex question that involves diagnosis and strategic decisions about what open or proprietary models to develop and for what purposes. China appears to be targeting a strategy of global AI-sharing that is empowered with open models. The July 2025 U.S. AI Action Plan highlighted the importance of the competition in open source,[27] and this competition is likely to become an even more important issue.

Improved public-private relationships and explicit partnerships, as well as recognizing the leading role of government in safeguarding national security and power, are essential. The public-private partnership dynamic appears across every approach. Every strategy relies heavily on some version of private-public synergy—whether through capital investment, public infrastructure, or safety enforcement—which reflects the United States' decentralized innovation ecosystem. However, the form of that relationship may be more cooperative or potentially antagonistic, depending on the strategic goals and approaches that the United States chooses to take.

Sequencing among strategy subcomponents is critical. Any comprehensive U.S. strategy will incorporate elements from all six archetypal strategies, and how the United States sequences the application of resources and influence will be essential. For example, favoring speed of development over safety and shared benefits might create crisis moments that ultimately slow medium-term progress. Successful sequencing will rely on AI's time frame: how such issues are approached over the next 24 months versus the next two decades, whether AI is set to take off to ASI or whether it will get there at all, what its broader effects will be, and whether their impact will be quick or more gradual over the next decades. U.S. strategy needs to hedge against either possibility.

[27] Executive Office of the President, 2025.

Comprehensive strategies should seek compounded returns. The goal of a combined or comprehensive strategy should be to find helpful feedback loops and compounding returns on investments, whether in the technological ecosystem that the United States is seeking to establish or through the political relations it fosters. The United States needs analysis that can identify the policies and investments that are most likely to trigger such cascades and that can understand the critical nodes that provide outsize benefits. Additionally, analysis should be aimed at determining the baseline degree of success that is necessary in any given step to be able to achieve the others.

Trends and policies outside the realm of AI will play critical roles in determining the effects of the AI transition on societies and on global balances of power. Across all strategies, non-AI elements (such as existing allied and partnership relations) shape and structure economic exchanges, military balances of power, and asymmetries that will shape the deployment and development of AI. Domestically, such issues as perceived socioeconomic inequality, corrupted information environments, perceived loss of control and autonomy, the growth and ossification of massive bureaucracies, and the effectiveness of governing institutions will continue to create social tensions that fracture U.S. politics, weaken the United States as a strategic actor, and undermine its inherent appeal. The United States will need a realistic perspective on how AI shapes and is shaped by these elements to successfully implement its strategy. Ultimately, the United States requires a comprehensive strategy for national power and renewal—one that encompasses its AI strategy but also addresses larger issues of society and economy.

A Comprehensive U.S. AI Strategy: Discovering the Foundation Stones

Addressing all these challenges and formulating a coherent, comprehensive U.S. AI strategy will not be easy, especially given the poor state of knowledge about AI throughout the U.S. government. Yet the urgency is clear: Dramatic advances in AI that have profound implications for U.S. interests could be available within the next 12 to 24 months. The U.S. government needs to start building the foundation of that strategy without knowing yet what the future AI landscape will look like. In other words, as we argued previously, U.S. AI strategy will be less a set document and more a continually evolving process, which adapts to new information while keeping fundamental goals and priorities in focus. Such flexible, adaptive strategy-making for AI will mean taking an array of actions to support the strategy's goals.[28] We can begin immediately to design strategic elements in two categories of actions: *no-regret* steps the U.S. government can take immediately that will position the United States for success in just about any scenario and the development of mechanisms for effective adaptive responsiveness over time. In the following sections, we examine each of these components, laying out the foundations of an adaptive U.S. AI strategic architecture.

What we outline below is not a comprehensive AI strategy. Rather, given the profound uncertainties and necessarily fluid nature of the U.S. response, we have tried to identify areas in which actions, processes, and mindsets can already be set up to create the essential foundation for an ongoing strategic response. Our recommendations are based on a set of conceptual understandings and include initial steps to get started.

Our analysis supports several important steps that the United States can take to set an effective foundation for its AI strategy that will be relevant under almost any scenario. These constitute the first and most-predictable component of an AI strategy. They remain fairly broad, but, in each case, we have offered an example of a specific policy initiative that would reflect such a move.

First, *the United States should begin setting up institutions at various levels of governance (federal, state, and local) to provide the architecture for responsive adaptation in strategy*. This new class of institutions would

[28] Michael Mankins, "In Uncertain Times, the Best Strategy Is Adaptability," *Harvard Business Review*, August 24, 2022.

be designed to develop accurate, cross-governmental assessments of the evolving AI landscape and develop detailed contingency plans for various possible developments. This effort will demand a dramatic boost in the level of AI expertise in the public sector, especially at the federal level but also in states and municipalities. It is especially critical to anticipate and manage the implications of a potential transition to ASI. The further along that transition AI gets, the more conscious intention will be required on the part of the U.S. government to deal with the radical implications of stunningly powerful AI agents. Such a transition does not necessarily imply greater government control or any specific regulatory actions. The U.S. government is likely to need even more significant AI policy capabilities to build on the existing AI Strategy Office in the White House for ongoing awareness, as well as assessment of emerging opportunities and risks and skilled personnel to oversee AI adoption throughout the government. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: An AI Strategy Office in All Federal Agencies and Governors' Offices

The federal government and state governments should create offices that are specifically designed to create and—of equal importance—continually update their AI strategies, including technology investments, workforce retraining, and loss-of-control preparation. These offices would not be focused on developing a single annual strategy document as much as an ongoing strategy process. They would fall under the general coordination of the top-level AI Strategy Office and would work with each other and closely engage with the AI Readiness Center and Open-Source Awareness Institute (suggested in the third practical initiative).

Second, *both public and private actors should begin developing contingency plans for several key classes of risk or opportunity events*,[29] such as loss of control, bioweapon development, or accelerating job-loss scenarios. Focusing in the first instance on high-impact, high-risk scenarios, these contingency plans should be developed while taking broader societal perspectives into account.[30] Regular exercises and games would be important to developing plans that are ready to be implemented at a moment's notice. Again, this imperative is connected to the prospects for vastly more-powerful AI, which will create new contingency risks related to the power of the AI itself, its autonomous capabilities, and ways it could be employed for harm. The potential for the emergence of some form of ASI over the next two years magnifies the importance of such contingency planning. It is the most essential baseline tool for preparing for such a transition. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: A National AI Readiness Center

This center would be responsible for identifying possible AI safety, alignment, and misuse issues, gaming out possible scenarios, and developing playbooks for response at both the federal and state levels. Conceptualized along the lines of the National Counterterrorism Center or other cross-government issue-based fusion cells, this center would engage far more significantly with the private and nonprofit sectors.

[29] Jim Mitre and Joel B. Predd, *Artificial General Intelligence's Five Hard National Security Problems*, RAND Corporation, PE-A3691-4, February 2025.

[30] Matan Chorev and Joel B. Predd, "America Should Assume the Worst About AI," *Foreign Affairs*, July 22, 2025.

Third, *to sustain an adaptive and flexible strategy, the United States will need the best possible information about the emerging AI transition*. The United States should pursue a program to develop better capabilities and information to track the capabilities of models in the U.S. ecosystem and internationally, monitor the intentions of other countries, and develop a better understanding of where competitors, such as China, stand. AI for both formal intelligence and open-source investments will also be increasingly important, as will the better integration of anticipatory intelligence, futures analysis,[31] and strategy-making into existing processes (which may require new offices and new specialties). Such measures should include up-to-date information about the status of models that are being developed in AI labs to have a clear sense of the future emerging model capabilities as AI models get closer to more powerful thresholds. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: A National Open-Source AI Awareness Institute

The federal government could create a small organization with technical experts, as well as general open-source intelligence analysts, to collect and assess all available information in the public domain to gauge the status of AI-model progress and associated infrastructure construction in the United States and abroad. The government could empower the institute with legal rights to access all proprietary information in the AI labs to ensure full transparency.

Fourth, *any plausible AI future is going to demand that the American people have a high degree of AI literacy*, including an understanding of how models work, their strengths and drawbacks, where they can be applied for great advantage, and what risks they pose. Such knowledge will be especially crucial among key groups in knowledge fields, such as educators, scientists, researchers, and medical professionals. Two absolutely central issues that are represented in the strategies—AI diffusion and application and empowering the American people—will depend significantly on the degree to which large numbers of Americans become empowered rather than alienated by the possibilities of AI tools. That is true for both employees performing specific tasks in established work environments and technology entrepreneurs trying to use AI tools to create new businesses or even entirely new industries. AI skills and capabilities will also be critical determinants of whether the technology enhances people's agency in their personal lives. Having a population that is capable of using the new technology will be more urgent and more important than during the early Industrial Revolution, when many emerging industries required relatively primitive labor and when the transition was played out over decades. It will be essential for state, local, and federal governments to promote actively a process of AI skill development. (For an example of a practical initiative based on this recommendation, see the following box.)

[31] J. Peter Scoblic, "Strategic Foresight in U.S. Agencies: An Analysis of Long-Term Anticipatory Thinking in the Federal Government," New America, last updated December 15, 2021.

Practical Initiative: Pilot Federal and State Programs in AI Application and Workforce Transition

Claude suggested the following version of such an initiative: "Comprehensive national program combining AI education in K–12, massive adult retraining initiatives, and 'AI co-pilot' tools for every American worker. Include stipends for participation and guarantees of AI-augmented employment placement."[a] ChatGPT added details, including "AI literacy and retraining initiatives across sectors, . . . support for labor-AI co-design, union partnerships, and human-centered AI interfaces, [and] programs for AI-in-government careers."[b]

[a] Claude, output from prompt by Michael Mazarr on the subject matter of the paragraph, version Claude Sonnet 4, August–September 2025.

[b] ChatGPT, output from prompt by Michael Mazarr on the subject matter of the paragraph, version ChatGPT 4o, August–September 2025.

Fifth, given the critical role of the practical application of AI models to real-world contexts in generating national AI advantage, *the United States will need to catalyze the diffusion and application of AI beyond model development*. Winning the race to create the best models is only part of the job. Prevailing in real-world applications is the only route to tangible competitive advantage, and there are reasons to believe China is positioning itself better for that challenge. Application is the route through which AI capability development produces tangible outcomes in the economy and society. Mastering the demands of AI diffusion and application will require significant efforts across the public, private, and nonprofit sectors. The goal is not to put government in the lead in making diffusion happen; rather, the U.S. government should encourage an adaptive, grassroots approach that generates thousands of experiments in using AI to improve society and power the economy. At all levels, governments can play an important role in making that happen, and such efforts need to be at the very core of any U.S. AI strategy. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: AI Diffusion Research Project Agency

The federal government could fund an office that is designed to undertake higher-risk AI applications to catalyze the diffusion of actual AI use throughout society. As with the Defense Advanced Research Projects Agency (DARPA), the resulting concepts and technologies would be available for general adoption. To serve other goals, a significant proportion of experiments could be directed to ensure that AI offers shared benefits and empowers autonomous human agency.

Sixth, the U.S. government should *accelerate AI safety initiatives* that increase the transparency of model capabilities, as well as the likelihood of detecting dangerous model characteristics before they are released. This imperative has been widely discussed, and there are many proposals to achieve these goals. Given the importance of safety to all the strategies, our analysis suggests that this element must be a priority even at the risk of delaying the rollout of new model capabilities. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: Trusted AI Alliance and Standards Initiative

Many experts and observers have suggested ways of enhancing AI safety. Claude proposed the idea of a "formal alliance structure with democratic partners for shared AI development standards, safety protocols, and mutual infrastructure access. Include technology transfer agreements, joint research funding, and coordinated approach to China." More ambitiously, it also nominated the possibility of an "AI Safety and Alignment Research Moonshot," with substantial funding for federal "AI alignment research, interpretability, and safety engineering. Create national labs dedicated to these challenges with mandatory participation from major AI companies. Include international research partnerships."[a] At a minimum, the United States should undertake what ChatGPT termed a "National AI safety research agency (analogous to DARPA or the Advanced Research Projects Agency for Health)."[b]

[a] Claude, output from prompt by Michael Mazarr on the subject matter of the paragraph, version Claude Sonnet 4, August–September 2025.

[b] ChatGPT, output from prompt by Michael Mazarr on the subject matter of the paragraph, version ChatGPT 4o, August–September 2025.

Seventh, *the United States should begin to lay the groundwork for a strategy to achieve and sustain a global AI network advantage*. A consistent theme across all the strategies we considered is the critical importance of network power in the integrated, AI-influenced digital future that is emerging. Given the potential for early network advantage and even potential lock-in for some components, the United States cannot afford to fall too far behind. U.S. AI technology leadership will contribute to this goal but will not assure it; conceivably, China could undercut U.S. providers on price or through assisting in AI's implementation. Although the market will determine much of this, state strategies make a difference: China is making a strong pitch to developing countries on digital networks and AI, and U.S. economic statecraft is simply not up to the challenge. Moreover, the market process itself will leave *AI deserts* in which infrastructure is too weak to support big applications. As we suggest in the following initiative, the risk of such outcomes can be partly offset by subsidizing access to U.S. AI capabilities. The United States could also offer risk-reduction guarantees for U.S. firms to undertake investment in targeted places and develop special strategies for partnership with regional hub nations. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: Global AI Development Corps and Open Tools Initiative

Claude suggested the idea of a "large-scale international development program deploying American AI expertise globally, combined with U.S.-developed open-source AI tools optimized for development applications. Include educational exchanges and capacity-building partnerships." This could include "[s]ubsidized global access to U.S.-developed open models and AI-for-good platforms (health, education, sustainability)."[a] Part of this program could be achieved with generous credits for compute and API [application programming interface] access. Presumably, any version of shared global AI development would also have to address infrastructure issues, such as where key nodes are to be built, by whom, and what access participating nations will have.

[a] Claude, output from prompt by Michael Mazarr on the subject matter of the paragraph, version Claude Sonnet 4, August–September 2025.

Eighth, f*ederal and state governments should undertake initial, experimental programs and policies that are designed to ensure shared benefits of AI among the American people*. Ensuring that AI produces shared benefits is critical to (1) safeguard against AI significantly contributing to worsening income and wealth inequalities; (2) promote AI use cases that improve general health, welfare, and economic security; and (3) encourage applications, norms, and regulations that promote rather than undermine autonomous human agency in the AI era. This is a very broad mandate that will require both a federal policy foundation and state and local tailoring based on the economy and its indicators. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: Sector- and Locality-Specific AI Benefit-Sharing Compact Experiments

Federal, state, and municipal governments can begin experiments in AI benefit-sharing of several types. A public-private consortium could take a sector in which the effects of AI jobs will come early—perhaps graphic design or software-coding—and undertake several different experiments in job retraining, task redesign, direct support, and other ways of easing the labor transition. Municipal governments could undertake experiments in how to bring the benefits of early AI application to lower-income citizens or depressed areas. States could create AI-driven innovation hubs in underserved communities and areas that are affected by deindustrialization.

Ninth, as a foundational element of securing U.S. global leadership, *the U.S. government should focus on attracting top talent both within its ranks and throughout the United States with the appropriate technical, administrative, and creative expertise*. Such efforts would be complementary to upskilling and workforce transition efforts, as well as early education AI literacy efforts. (For an example of a practical initiative based on this recommendation, see the following box.)

Practical Initiative: AI Fellowship Program

U.S. government departments and agencies have long used various forms of fellowship programs to attract top talent. Some of them are yearlong paid internships that expose the fellows to various offices within a department or across several departments. Other programs offer tuition support for longer work commitments. Some programs are focused on specific skills, such as the CyberForce Scholarship for Service initiative. The U.S. government could adopt a similar model for an AI fellowship that centers on technical, policy, and other AI-related expertise to constitute the next generation of AI-native public policy servants for the duration considered necessary to uptake their expertise and skills.

As suggested previously, one idea of any comprehensive AI strategy is to achieve compounding and mutually reinforcing advantages. Each element of a strategy should be designed to nest with and support the others. Achieving such compounding effects should be a major focus of the AI strategy offices. Claude phrased this goal in the following way:

> These [multiple] policies create a virtuous cycle: [I]nfrastructure enables deployment, workforce development ensures adoption, international partnerships create markets and legitimacy, safety research reduces

> barriers to rapid scaling, and global engagement builds the network effects that sustain long-term dominance. Critically, each investment makes the others more effective while creating irreversible advantages—once the U.S. becomes the essential infrastructure provider, safety standard-setter, and capacity-builder, maintaining that position becomes much easier than achieving it initially.[32]

Note that how China and others react to a given strategy is a critical factor. Although we have taken international response into account in this paper, our aim is to generate an initial view of strategic concepts. We will test reactions and second-order effects in future workshops and games.

Conclusion

Those who experienced the early phases of what would become the Industrial Revolution had no way of grasping the magnitude of change ahead of them: economically (in the transition from agrarian economies to industrial ones and the rise of working and middle classes), socially (in the redistribution of wealth and authority), technologically (in the emergence of new sources of energy, transportation, and other inventions), geopolitically (in the rise of the UK as a center of global power), and culturally (in the early signs of an interconnected global community). The scope and character of that epochal transformation became clear only in retrospect.

The United States has a chance to look across the horizon and wrestle with the possibilities of how AGI may transform our societies *before* it does so. The possibilities are both perilous and unshackling. At one end of the spectrum, there could be a total transformation of how our societies trade, fight, and live; however, it is more likely that there will enter a period of tumultuous but recognizable transition in which the United States will make choices that bear long-lasting consequences. No single strategy can navigate the wide variety of ways this change might occur. But, when confronted with such a dramatic paradigm shift, a good strategy is to address the most-essential requirements for success and provide a foundation for the flourishing of human ingenuity and resilience. That is our task now in developing a U.S. strategy for the AI era.

---

[32] Claude, output from prompt by Michael Mazarr on the subject matter of the paragraph, version Claude Sonnet 4, August–September 2025.

Abbreviations

AI	artificial intelligence
AGI	artificial general intelligence
ASI	artificial superintelligence
DARPA	Defense Advanced Research Projects Agency
EU	European Union
LLM	large language model
R&D	research and development
ROK	Republic of Korea
UK	United Kingdom

References

Aschenbrenner, Leopold, "Situational Awareness: The Decade Ahead," *Situational Awareness* blog, June 2024.

Bateman, Jon, Dan Baer, Stephanie A. Bell, Glenn O. Brown, Mariano-Florentino (Tino) Cuéllar, Deep Ganguli, Peter Henderson, Brodi Kotila, Larry Lessig, Nicklas Berild Lundblad, Janet Napolitano, Deborah Raji, Elizabeth Seger, Matt Sheehan, Aviya Skowron, Irene Solaiman, Helen Toner, and Polina Zvyagina, *Beyond Open vs. Closed: Emerging Consensus and Key Questions for Foundation AI Model Governance*, Carnegie Endowment for International Peace, July 2024.

Biden, Joseph R., "On Advancing the United States' Leadership in Artificial Intelligence; Harnessing Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence," memorandum, White House, October 24, 2024.

Chorev, Matan, and Joel Predd, "America Should Assume the Worst About AI," *Foreign Affairs*, July 22, 2025.

Department for Science, Innovation & Technology, *AI Opportunities Action Plan: Government Response*, Government of the United Kingdom, 2025.

Ding, Jeffrey, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, Princeton University Press, 2024.

Ding, Jeffrey, *Running the Right Artificial Intelligence Race: A National Strategy for AI Diffusion*, RAND Corporation, PE-A4165-1, September 2025. As of October 1, 2025: https://www.rand.org/pubs/perspectives/PEA4165-1.html

Egan, Janet, "The Way Forward: AI Diplomacy with Compute Leadership," in *Global Compute and National Security: Strengthening American AI Leadership Through Proactive Partnerships*, Center for a New American Security, July 2025.

European Commission, "Communication from the Commission: Artificial Intelligence for Europe," COM(2018) 237 final, April 25, 2018.

European Commission, "The AI Continent Action Plan," webpage, April 9, 2025. As of September 15, 2025: https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan

Executive Office of the President, *Winning the Race: America's AI Action Plan*, White House, July 2025.

Executive Order 14110, "Advancing United States Leadership in Artificial Intelligence Infrastructure," Executive Office of the President, January 14, 2025.

Government of Canada, "Pan-Canadian Artificial Intelligence Strategy," webpage, last updated December 3, 2024. As of September 15, 2025: https://ised-isde.canada.ca/site/ai-strategy/en

Government of the United Kingdom, "The UK's National AI Strategy," flow diagram, undated.

Government of the United Kingdom, *National AI Strategy*, version 1.2, September 2021.

Hendrycks, Dan, Eric Schmidt, and Alexandr Wang, "Superintelligence Strategy: Expert Version," arXiv, arXiv:2503.05628v2, April 14, 2025.

Jin, Berber, and Joanna Stern, "Anthropic CEO Says AI Could Surpass Human Intelligence by 2027," *Wall Street Journal*, January 21, 2025.

Kantrowitz, Alex, "Google DeepMind CEO Demis Hassabis: The Path to AGI, LLM Creativity, and Google Smart Glasses," *Big Technology* blog, January 23, 2025. As of September 15, 2025: https://www.bigtechnology.com/p/google-deepmind-ceo-demis-hassabis

Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, "AI 2027," AI Futures Project, April 3, 2025.

Mankins, Michael, "In Uncertain Times, the Best Strategy Is Adaptability," *Harvard Business Review*, August 24, 2022.

Marcellino, William, Lav Varshney, Anton Shenk, Nicolas M. Robles, and Benjamin Boudreaux, *Charting Multiple Courses to Artificial General Intelligence*, RAND Corporation, PE-A3691-1, April 2025. As of September 15, 2025:
https://www.rand.org/pubs/perspectives/PEA3691-1.html

Mazarr, Michael J., *The Societal Foundations of National Competitiveness*, RAND Corporation, RR-A499-1, 2022. As of September 15, 2025:
https://www.rand.org/pubs/research_reports/RRA499-1.html

Mazarr, Michael, "AI and the Challenge of Adaptive Strategy," *Geopolitics of AGI*, Substack, July 2, 2025. As of September 15, 2025:
https://geopoliticsagi.substack.com/p/ai-and-the-challenge-of-adaptive

Mitre, Jim, and Joel B. Predd, *Artificial General Intelligence's Five Hard National Security Problems*, RAND Corporation, PE-A3691-4, February 2025. As of September 15, 2025:
https://www.rand.org/pubs/perspectives/PEA3691-4.html

Morris, Meredith Ringel, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg, "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv, arXiv:2311.02462v4, June 5, 2024.

Narayanan, Arvind, and Sayash Kapoor, "AI as Normal Technology," *AI as Normal Technology* blog, April 15, 2025. As of September 15, 2025:
https://www.normaltech.ai/p/ai-as-normal-technology

Özcan, Bengüsu, Daan Juijn, Jakob Graabak, and Sam Bogerd, "Advanced AI: Possible Futures—Five Scenarios for How the AI-Transition Could Unfold," Centre for Future Generations, July 10, 2025. As of September 15, 2025:
https://cfg.eu/advanced-ai-possible-futures/

Plumb, Radha Iyengar, and Michael C. Horowitz, "What America Gets Wrong About the AI Race," *Foreign Affairs*, April 18, 2025.

Radu, Roxana, "Steering the Governance of Artificial Intelligence: National Strategies in Perspective," *Policy and Society*, Vol. 40, No. 2, June 2021.

Scoblic, J. Peter, "Strategic Foresight in U.S. Agencies: An Analysis of Long-Term Anticipatory Thinking in the Federal Government," New America, last updated December 15, 2021.

Stanford University Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2024*, Stanford University, 2024.

Tyrangiel, Josh, "Sam Altman on ChatGPT's First Two Years, Elon Musk and AI Under Trump," Bloomberg, January 5, 2025.

Webster, Graham, Rogier Creemers, Elsa Kania, and Paul Triolo, "Full Translation: China's 'New Generation Artificial Intelligence Development Plan' (2017)," DigiChina, Stanford University, August 1, 2017.

About the Authors

Michael Mazarr is a senior political scientist at RAND. His primary interests are U.S. defense policy and force structure, disinformation and information manipulation, East Asian security, nuclear weapons and deterrence, and judgment and decisionmaking under uncertainty. He holds a Ph.D. in public affairs.

Anca Agachi is a policy analyst at RAND. Her research interests include European security and defense, alliance dynamics, global governance, the security implications of emerging technologies, and strategic foresight. She holds an M.A. in international affairs.