



**Before the
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
Washington, D.C. 20500**

**COMMENTS OF THE
COMPETITIVE ENTERPRISE INSTITUTE**

In the Matter of)	
)	
Office of Science and Technology)	87 FR 5876
Policy)	Document Number 2022-02161
)	
RFI Response:)	
National Artificial Intelligence)	
Research and Development)	
Strategic Plan)	
 March 4, 2022		 Ryan Nabil COMPETITIVE ENTERPRISE INSTITUTE 1310 L Street NW, 7 th Floor Washington, DC 20005 (202) 331-1010

Introduction

On behalf of the Competitive Enterprise Institute (CEI), I respectfully submit these comments in response to the Office of Science and Technology Policy's request for comments on updating the National Artificial Intelligence Research and Development Strategic Plan.¹ Founded in 1984, the Competitive Enterprise Institute is a non-profit research and advocacy organization that focuses on regulatory policy from a pro-market perspective. CEI works to promote policies that can help boost technological innovation and American leadership in areas such as artificial intelligence and machine learning and emerging technologies that depend on such innovation.

The Competitive Enterprise Institute appreciates the recognition by both the Trump and Biden Administrations of the need to create a more favorable regulatory environment in which artificial intelligence (AI) and AI-enabled emerging technologies can thrive and promote American economic growth and competitiveness. To that end, CEI recognizes the increasingly important role played by the Office of Science and Technology Policy (OSTP), the National Science and Technology Council (NSTC)'s Select Committee on Artificial Intelligence, the NSTC Machine Learning and AI Subcommittee, and the National AI Initiative Office. As these regulatory bodies seek to update the National AI Research and Development (R&D) Strategic Plan, they have an opportunity to strengthen America's position as a global center of AI innovation. To accomplish that goal, the OSTP and the NSTC need to make several updates to the current AI strategy.

Specifically, the National AI R&D Strategic Plan would benefit from updates in five areas.

- First, while the strategic plan recognizes the essential role of the private sector in promoting AI innovation, it needs to provide more concrete steps to engage the private sector in AI research and development projects.
- Second, to ensure that taxpayer dollars are utilized effectively, the AI strategic plan should propose a framework to track and evaluate the effectiveness of R&D expenditure and grants to various recipients in different AI subdisciplines.
- Third, the strategic plan would benefit from a more nuanced understanding of the AI R&D and regulatory approaches in other countries. For example, reviewing AI policies of other nations could help inform why the U.S. government should allocate a higher share of research spending to multidisciplinary AI projects and prioritize accuracy over algorithmic transparency as a goal for AI systems.
- Fourth, while the AI strategy recommends developing shared AI datasets for academic and private sector use, more details are needed on such proposals.
- Fifth, the National AI R&D Strategic Plan should propose a federal AI sandbox program to incentivize the private sector to play a more important role in AI innovation. By allowing private companies and research institutions to test innovative AI systems for a limited time, such a program can help promote technological innovation, enhance regulatory understanding of AI, and help craft market-friendly regulatory frameworks and technical standards for AI systems.

¹ Office of Science and Technology Policy (OSTP), "Request for Information to the Update of the National Artificial Intelligence Research and Development Strategic Plan," *Federal Register*, Vol. 87, No. 22 (February 2, 2022), <https://www.federalregister.gov/documents/2022/02/02/2022-02161/request-for-information-to-the-update-of-the-national-artificial-intelligence-research-and>.

I. The National AI R&D Strategic Plan Needs to Better Engage the Private Sector

The private sector and academic institutions play a crucial role in the development of AI technologies.² Given that reality, a successful AI strategy needs to closely engage technology companies, startups, and research institutions. The 2019 National AI R&D Strategic Plan recognizes the private sector’s essential role in promoting artificial intelligence and provides several recommendations to enhance collaboration with the private sector. For instance, it proposes creating joint public-private collaboration, increasing the availability of public datasets, and expanding AI training and fellowship opportunities to meet workforce R&D needs.³ Despite such proposals, the strategy would benefit from a greater emphasis on engaging the private sector and academic institutions in promoting AI innovation, for example, by creating a federal artificial intelligence regulatory sandbox program, as discussed later in this comment.

II. Developing Mechanisms to Track and Evaluate Artificial Intelligence R&D Spending

The National AI R&D Strategic Plan should propose a framework to better track the allocation and impact of AI-related research and development projects across federal agencies, research institutions, companies, and other recipients of federal AI R&D grants. Despite the growing federal expenditure on AI-related research and development activities, there appears to be a scarcity of efforts in tracking how this money is spent and how it impacts AI innovation.⁴

Greater transparency and more precise information about federal AI expenditure and its impact on innovation within different AI subdisciplines can help policymakers allocate R&D resources more effectively. For example, have resources allocated to specific AI subdisciplines—such as computer vision—led to demonstrably better research outcomes than in other areas? Are certain federal agencies and academic institutions more effective at utilizing research grants than others? Collecting and analyzing data to answer these questions can significantly improve policymakers’ ability to make evidence-based R&D spending decisions.

Some long-term AI research projects will require several years before R&D efforts show results—especially in “general AI” and areas of machine learning research that do not appear to have immediate commercial applications.⁵ However, tracking spending can nonetheless help compare the effectiveness of similar short- and long-term projects by different agencies, research institutions, and companies. That could not only help U.S. policymakers allocate more resources to more promising AI subdisciplines, but it might also help improve competition between different recipients of federal research grants.

² White House, Office of the President, Select Committee on Artificial Intelligence of the National Science and Technology Council (NSTC), *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, June 2019, p. iii, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>; Lauren A. Kahn, “US Leadership in Artificial Intelligence Is Still Possible,” Council on Foreign Relations, October 28, 2021, <https://www.cfr.org/blog/us-leadership-artificial-intelligence-still-possible>.

³ NSTC, *The 2019 National AI R&D Strategic Plan*, p. iii.

⁴ Jon Harper, “Federal AI Spending to Top \$6 Billion,” *National Defense*, February 10, 2021, [https://www.nationaldefensemagazine.org/articles/2021/2/10/federal-ai-spending-to-top-\\$6-billion](https://www.nationaldefensemagazine.org/articles/2021/2/10/federal-ai-spending-to-top-$6-billion).

⁵ The National AI R&D Strategic Plan defines general AI as the type of AI intended to “create systems that exhibit the flexibility and versatility of human intelligence in a broad range of cognitive domains, including learning, language, perception, reasoning, creativity, and planning” (NSTC, *The 2019 AI R&D Strategic Plan*, pp. 10–11).

III. Lessons from the Successes and Shortcomings of Other Countries' AI Policies

The National AI Strategic R&D Plan would benefit from closer examination of how other countries allocate research spending, their regulatory approach toward AI, and the extent to which these policies have been successful. Due to the potential military applications of many AI-enabled technologies, other countries' AI strategies—particularly those of adversarial nations—are often viewed as a threat to America's national security and technological competitiveness.

However, AI policies and developments in other countries also provide the opportunity to better understand which R&D and regulatory approaches have been successful elsewhere. Policymakers should exercise caution in making such comparisons, as the regulatory experience from other jurisdictions might have limited applicability to the United States. However, awareness of those broader trends can help the U.S. capitalize on the successes of different countries and avoid their regulatory mistakes. To maximize the benefit of this comparative approach, the strategic plan could propose mechanisms to conduct annual reviews of the global AI research and regulatory landscape and an evaluation of its successes and failures.

For example, the National AI R&D Strategic Plan could benefit from a closer examination of European and Chinese approaches toward interdisciplinary AI research. While the European Union's restrictive approach to AI risks harming innovation in certain sectors,⁶ several European countries have designed innovative AI R&D strategies at the national level. For example, the French government has proposed creating a network of interdisciplinary institutions to promote high-level AI research in multiple disciplines. Outside of the European Union, the British government has launched initiatives to bolster multidisciplinary artificial intelligence research and foster AI-enabled innovation in insurance and legal services.⁷ In China—one of the two leading sources of AI innovation alongside the United States—the State Council has created similar initiatives to encourage cross-disciplinary academic research at the intersection of artificial intelligence, economics, psychology, and other core disciplines.⁸

Studying and evaluating these countries' approaches might provide American policymakers insights into the extent to which existing R&D resources should be devoted to interdisciplinary AI projects. To that end, the U.S. government could also create an AI sandbox and other innovative programs to capitalize on the expertise of academic institutions and the private sector and promote cross-disciplinary AI research in fields such as finance, medicine, and physics.

⁶ Ryan Nabil, "The EU's Recently Proposed Artificial Intelligence Act Goes Too Far," *The National Interest*, August 16, 2021, <https://nationalinterest.org/blog/buzz/eu-s-recently-proposed-artificial-intelligence-act%2%A0goes-too-far-191733>; Benjamin Muller, *How Much Will the Artificial Intelligence Act Cost Europe?* (Brussels: Center for Data Innovation, July 2021), <https://www2.datainnovation.org/2021-aia-costs.pdf>.

⁷ Castro and New, "2016 National AI R&D Strategic Plan," p. 3.

⁸ State Council of China, "国务院关于印发新一代人工智能发展规划的通知" ["Notice of the State Council on the Release of the New Generation Artificial Intelligence Development Plan"], July 20, 2017, http://www.gov.cn/zhengce/content/2017-07/20/content_5211996.htm; Castro and New, "2016 National AI R&D Strategic Plan," p. 3. For an English translation of the Chinese government's 2017 AI Development Plan, see State Council of China, "China's 'New Generation Artificial Intelligence Development Plan'" [translated by Graham Webster, Roger Creemers, Paul Triolo, and Elsa Kania], *New America*, August 1, 2017, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.

Along with successful innovative strategies of other countries, it is equally important to examine potential regulatory mistakes. For instance, despite objections from government bodies and non-profit research organizations,⁹ the UK and the EU have sought to promote algorithmic transparency as a major objective for AI technologies.¹⁰ Likewise, many American policymakers and experts increasingly advocate transparency as a goal in developing AI systems.¹¹ In attempting to understand AI systems better, the National AI R&D Strategic Plan has proposed “improving fairness, transparency, and accountability-by-design” as objectives for government-funded AI research projects.¹²

While mechanisms to improve fairness and reduce biases are worthwhile goals, mandating high levels of transparency could detract from designing effective and accurate AI programs.¹³ That is all the more likely if the same level of transparency requirements were to be applied to AI systems across all industries. For example, high levels of algorithmic transparency are crucial to prevent government abuses of power and civil liberties violations if AI systems are used in fields such as criminal justice and law enforcement. However, transparency concerns are much less important in areas like cybersecurity and medical diagnosis—where the accuracy of preventing cyberattacks and detecting diseases is much more important than the explainability or transparency of underlying algorithms. Therefore, the National AI R&D Strategic Plan should not prioritize transparency over accuracy as a research objective—except in limited and specified cases such as AI systems for criminal justice, law enforcement, and human resources.¹⁴

Since AI systems typically involve complex neural networks and algorithms that produce effective outcomes, their underlying processes are not transparent even to the programmers that design such systems.¹⁵ For instance, medical researchers at New York’s Mount Sinai Hospital developed Deep Patient, an AI program that can predict whether a patient has contracted a specific disease. The AI system, which is reported to be substantially better than any comparable systems, trained on the medical data of 700,000 patients across several hundred variables. Deep Patient can accurately provide medical diagnostics, but its designers cannot accurately describe

⁹ UK Government Office for Science, *Artificial Intelligence: Opportunities and Implications for the Future of Decision Making* (London: Government Office for Science, February 12, 2016), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/566075/gs-16-19-artificial-intelligence-ai-report.pdf; Nick Wallace and Daniel Castro, “The Impact of the EU’s New Data Protection Regulation on AI,” Center for Data Innovation, March 27, 2018, p. 4, <https://www2.datainnovation.org/2018-impact-gdpr-ai.pdf>.

¹⁰ UK Government, Central Digital and Data Office, “Algorithmic Transparency Standard,” November 29, 2021, www.gov.uk/government/collections/algorithmic-transparency-standard; Tambiama Madiega, “EU Guidelines on Ethics in Artificial Intelligence: Context and Implementation,” European Parliament Briefing PE 650.163 (September 2019), p. 4, [https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/640163/EPRS_BRI\(2019\)640163_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/640163/EPRS_BRI(2019)640163_EN.pdf).

¹¹ For example, see Electronic Privacy Information Center, “State Artificial Intelligence Policy,” n.d., accessed March 3, 2022, <https://epic.org/state-artificial-intelligence-policy/>; Artificial Intelligence Capabilities and Transparency Act of 2021, S. 1705, 117th Cong. (2017), <https://www.congress.gov/bill/117th-congress/senate-bill/1705/text?r=82&s=1>.

¹² NSTC, *The 2019 National AI R&D Strategic Plan*, p. 21.

¹³ Daniel Castro and Joshua New, “Regulatory Comments in Response to the Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan,” Center for Data Innovation, October 26, 2018, pp. 4–5, <https://www2.datainnovation.org/2018-nitrd-ai-r%26d.pdf>.

¹⁴ *Ibid.*

¹⁵ *Ibid.*, p. 4.

how the program arrives at such a diagnosis.¹⁶ Mandating high levels of algorithmic transparency from such programs can detract from cutting-edge AI innovation in fields ranging from healthcare to nuclear science.

Nevertheless, to bolster leadership in shaping global AI norms, Britain and the EU are increasingly advocating algorithmic transparency as a main AI policy goal.¹⁷ Although fixating on transparency can harm the development of innovative yet unexplainable AI systems, potential benefits remain limited. As the UK Government Office for Science explained in a 2016 report:

Most fundamentally, transparency may not provide the proof sought: Simply sharing static code provides no assurance it was actually used in a particular decision, or that it behaves in the wild in the way its programmers expect on a given dataset.¹⁸

Studying these international regulatory developments and criticisms can help American policymakers avoid potential mistakes like mandating transparency as an overarching goal for all types of AI systems. By creating processes to review international regulatory developments, the National AI R&D Strategic Plan can help U.S. policymakers design an innovation-friendly approach toward artificial intelligence.

IV. Encouraging Academic and Private Sector Innovation by Providing Access to Non-Sensitive Government Database

The lack of access to data remains a significant challenge for the development of novel AI technologies, especially for startups and businesses without the resources of Big Tech companies. Creating innovative AI systems requires high-quality datasets on which AI systems can be trained. The costs associated with creating, cleaning, and preparing such datasets for training AI systems remain too high for many businesses and academic institutions.¹⁹ For example, Google London-based subsidiary DeepMind's AlphaGo software made headlines in March 2016 when it defeated the human champion of the Chinese game Go.²⁰ The cost to train datasets for building this program was more than \$25 million in hardware alone.²¹

¹⁶ Castro and New, p.4; Will Knight, "The Dark Secret at the Heart of AI," *MIT Technology Review*, April 11, 2017, <https://www.technologyreview.com/s/604087/the-darksecret-at-the-heart-of-ai/>.

¹⁷ UK Government, Central Digital and Data Office, "UK Government Publishes Pioneering Standard for Algorithmic Transparency," news release, November 29, 2021, <https://www.gov.uk/government/news/uk-government-publishes-pioneering-standard-for-algorithmic-transparency>; Madiega, "EU Guidelines on AI Ethics."

¹⁸ UK Government, Office for Science, *Artificial Intelligence: Opportunities and Implications for the Future of Decision Making* (London: UK Government Office for Science, November 9, 2016), <https://www.gov.uk/government/publications/artificial-intelligence-an-overview-for-policy-makers>; Castro and New, pp. 4–5.

¹⁹ NSTC, *The 2019 National AI R&D Strategic Plan*, p. 28; Daniel E. Ho, Jennifer King, Russell C. Wald, and Christopher Wan, *Building a National AI Research Resource; A Blueprint for the National Research Cloud* (Stanford: Stanford University Human-Center Artificial Intelligence and Stanford Law School, October 2021), pp. 35–41, https://hai.stanford.edu/sites/default/files/2022-01/HAI_NRCR_v17.pdf.

²⁰ DeepMind. "The Story of AlphaGo So Far," n.d., accessed February 28, 2022, <https://deepmind.com/research/alphago/>; Jeffrey Ding, "Deciphering China's AI Dream," Centre for the Governance of AI, Future of Humanity Institute, University of Oxford, March 2018, p. 7. https://www.fhi.ox.ac.uk/wp-content/uploads/Deciphering_Chinas_AI-Dream.pdf.

²¹ Elizabeth Gibney, "Self-taught AI is best yet at strategy game Go," *Nature*, October 18, 2017, <https://www.nature.com/articles/nature.2017.22858>.

Recognizing this challenge, the National AI R&D Strategic Plan recommended the development of shared datasets that startups, businesses, and research institutions can use to create and train AI programs.²² But despite this commitment, progress in this area appears to be slow. That is why the AI strategic plan needs to outline more concrete steps to publish high-quality datasets using the vast amount of non-sensitive and non-personally identifiable data already at the federal government’s disposal. Under the 1974 Privacy Act, U.S. government agencies have not created a central repository of data, which is important because of the privacy and cybersecurity risks that a central data repository of sensitive information would face.²³

However, different U.S. agencies also have access to a wide range of non-personally identifiable and non-sensitive datasets intended for public use—such as the National Oceanic and Atmospheric Administration’s climate data and the National Aeronautics and Space Administration (NASA)’s non-confidential space-related data.²⁴ Making such data readily available to the public can allow AI innovation in weather forecasting, transportation, astronomy, and other underexplored subjects.²⁵ Therefore, the AI strategy should propose a framework that enables the OSTP and the NSTC to work with government agencies and ensure that non-sensitive *and* non-personally identifiable data—intended for public use—are made available in a format suitable for AI research by the private sector and research institutions.

The OSTP, the NSTC, and the National AI Initiative Office could use the federal government’s FedRAMP data classification as a general framework to develop a strategy for which data should be included in public datasets. The FedRAMP framework divides government-stored data into three distinct types:

- 1) Low-impact risk data meant for public use;
- 2) Moderate-impact risk data, which are controlled, unclassified data (e.g., personally identifiable information) unavailable to the public; and
- 3) High-impact risk data, which contain “sensitive federal information,” such as law enforcement and emergency services information.²⁶

To minimize privacy and cybersecurity risks, the OSTP and the NSTC should propose that the datasets only contain low-impact risk data intended for public use. The OSTP, the NSTC, and other relevant regulatory authorities should ensure appropriate data controls and privacy standards so that these datasets do not erroneously include sensitive information, uphold cybersecurity best practices, and are provided in a format suitable for training AI systems.

V. Creating a Federal Artificial Intelligence Regulatory Sandbox Program

To develop the regulatory understanding of emerging AI technologies and better engage the private sector, the National AI R&D Strategic Plan should propose the creation of a federal AI sandbox program. Given the rapid evolution of AI-enabled technologies, there is a growing need

²² NSTC, *The 2019 National AI R&D Strategic Plan*, p. 28.

²³ Privacy Act of 1974, 5 U.S.C. § 552a (2012); Ho et al., *National AI Research Resource*, p. 35.

²⁴ Ho et al., *National AI Research Resource*, p. 47.

²⁵ *Ibid.*

²⁶ *Ibid.* p. 39.

to better understand their ethical, legal, and societal implications.²⁷ For example, as the strategic plan notes, “the unusual complexity and evolving nature” of AI systems mean that “ensuring the safety of AI systems” remains a challenge.²⁸ Therefore, understanding such systems would constitute an important first step toward crafting safe, market-friendly regulatory frameworks and technical standards for AI systems and AI-enabled technologies.²⁹

At the same time, there is a growing need to engage the private sector in promoting AI innovation. These objectives could benefit from a policy tool commonly used to advance financial technology (FinTech) called “regulatory sandboxes.”

“Regulatory sandbox” programs provide companies with an experimenting space that allows them to offer innovative products and services under a frequently lightened regulatory framework for a limited period.³⁰ The United Kingdom’s Financial Conduct Authority created the world’s first FinTech sandbox program in 2015. Since then, regulators in Australia, Hong Kong, Singapore, and other innovative jurisdictions have launched similar programs.³¹ The federal Consumer Financial Protection Bureau and more than 10 U.S. states have launched sandbox programs to promote technological innovation in finance and insurance.³² In addition to financial services, the Utah Supreme Court has also created a regulatory sandbox program that allows non-legal firms to provide certain innovative legal services.³³

Although FinTech sandbox programs are becoming increasingly common, AI sandbox programs remain a largely underexplored idea.³⁴ In its 2021 Artificial Intelligence Act, the European Commission proposed an AI sandbox program.³⁵ However, its success will depend on its regulatory design and the extent to which the sandbox prioritizes technological innovation and regulatory learning. If implemented correctly, an AI regulatory sandbox program can be significantly helpful in promoting U.S. AI innovation. To that end, the National AI R&D Strategic Plan could recommend a sandbox program that will accept participants based on

²⁷ NSTC, *The 2019 National AI R&D Strategic Plan*, pp. 19–22.

²⁸ *Ibid.*, p. 5.

²⁹ *Ibid.*, pp. 14–18, 23–26.

³⁰ Ryan Nabil, “How Regulatory ‘Sandboxes’ Can Boost U.S. Technological Innovation,” *Real Clear Markets*, August 12, 2021, https://www.realclearmarkets.com/articles/2021/08/12/how_regulatory_sandboxes_can_boost_us_technological_innovation_789620.html

³¹ Financial Conduct Authority, “Financial Conduct Authority’s regulatory sandbox opens to applications,” May 9, 2016, <https://www.fca.org.uk/news/press-releases/financial-conduct-authority’s-regulatory-sandbox-opens-applications>.

³² World Bank, *Global Experiences from Regulatory Sandboxes* (Washington, DC: World Bank Fintech Note No. 8, 2020), <https://openknowledge.worldbank.org/handle/10986/34789>; Nabil, “How Regulatory ‘Sandboxes’ Can Boost U.S. Technological Innovation.”

³³ Utah Supreme Court, Office of Legal Services Innovation, “January 2022 Activity Report,” n.d., accessed February 27, 2022, <https://utahinnovationoffice.org>; Ryan Nabil, “Regulatory sandbox programs can promote legal innovation and improve access to justice,” *The Hill*, October 9, 2021, <https://thehill.com/opinion/judiciary/576041-regulatory-sandbox-programs-can-promote-legal-innovation-and-improve-access>.

³⁴ Aljoscha Burchardt, “Steckt die KI in den Sandkasten!” [“Put AI in a Sandbox!”], *Die Zeit*, December 10, 2020, <https://www.zeit.de/2020/52/kuenstliche-intelligenz-suchmaschinen-navigationen-training-forschung>.

³⁵ European Commission, “Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts,” Proposal for a Regulation of the European Parliament and of the Council, COM (2021) 206 Final, April 21, 2021, <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1623335154975&uri=CELEX%3A52021PC0206>.

proposed AI innovation and its potential benefit to developing novel AI technologies.³⁶ Such a program would not only benefit consumers and the participating companies; it would also enable regulators to gain firsthand insights into AI technologies and help craft market-friendly regulatory frameworks and technical standards.

Regulators could also create sandbox programs to target innovation in specific areas—such as human-machine interaction and probabilistic reasoning—which the strategic plan identifies as areas in need of further research.³⁷ For example, current AI systems are ill-equipped to translate heavily accented speech or speech in a noisy environment.³⁸ A thematic sandbox program targeting natural language processing could incentivize companies and researchers to offer innovative AI-enabled products in this area. Likewise, a thematic sandbox aimed at developing AI-enabled cybersecurity and encryption tools can help encourage market innovations to counter growing cybersecurity challenges.³⁹ Sandbox participants testing innovative products in these areas can go a long way towards helping lawmakers and regulators better understand these emerging areas of artificial intelligence and develop innovation-friendly regulatory frameworks accordingly.

The AI regulatory sandbox concept remains novel. Therefore, designing an effective program will require creative thinking from the Office of Science and Technology Policy, the National Science and Technology Council, and the National AI Initiative Office. For example, regulators will need to define the type of AI systems and AI-enabled products and services eligible for participating in the sandbox. Given the current technological limitations, an AI sandbox might need to be restricted to 1) “limited AI” systems that perform tasks in specific and well-defined domains like speech recognition, translation, and medical diagnosis and 2) projects where measurable technological advances are possible within the typical sandbox testing period of one to two years.⁴⁰

Furthermore, the systems, products, and services eligible for the sandbox could fall under the jurisdictions of multiple regulators. Such a development might ultimately require a legal framework that defines the supervisory role of different regulators in operating the AI sandbox program in cases of overlapping jurisdiction.

Designing an effective AI sandbox will also require modifications to the existing FinTech sandbox models. For example, unlike in financial services, academic institutions remain a leading source of AI innovation.⁴¹ Typically, FinTech sandbox programs provide participants

³⁶ To better understand the regulatory designs of major fintech sandbox programs in the United States and selected foreign jurisdictions, and how they might inform the regulatory design of an AI sandbox program, readers are advised to consult the author’s upcoming CEI report on regulatory sandbox programs.

³⁷ NSTC, *The 2019 National AI R&D Strategic Plan*, p. 12.

³⁸ *Ibid.*

³⁹ Arthur Herman, “The Executive’s Guide to Quantum Computing and Quantum-secure Cybersecurity,” Hudson Institute, April 3, 2019, <https://www.hudson.org/research/14930-the-executive-s-guide-to-quantum-computing-and-quantum-secure-cybersecurity>.

⁴⁰ In contrast to limited AI, the goal of “general AI” is to “create systems that exhibit the flexibility and versatility of human intelligence in a broad range of cognitive domains, including learning, language, perception, reasoning, creativity, and planning” (NSTC, *The 2019 National AI R&D Strategic Plan*, pp. 10–11).

⁴¹ NSTC, *The 2019 National AI R&D Strategic Plan*, p. iii.

with exemptive regulatory relief and waived or expedited registration processes.⁴² However, many research institutions might not seek to commercialize the algorithms and technologies they test in a sandbox. Therefore, exemptive regulatory relief might be of limited benefit to such institutions. Furthermore, for many research institutions and companies, the lack of access to high-quality datasets and cloud-based computing power might pose a greater obstacle than regulatory barriers.⁴³ Thus, a way to incentivize participation in an AI sandbox program could be to grant limited cloud-based computing power for the duration of the sandbox test.⁴⁴

Conclusion

The Office of Science and Technology Policy, the National Science and Technology Council, and the National AI Initiative Office need to adopt a realistic approach, objectives, and scope for a national AI R&D plan. AI's general-purpose nature—combined with its diffusion across many sectors and the rapidly changing technological developments—limits the extent to which a national strategy can significantly improve AI innovation across the economy. Recognizing this challenge, the Biden administration and administrative agencies should focus on enabling a wide range of actors, from tech startups to academic and financial institutions, to play a role in promoting American AI innovation.

Given the rapid change and growth of AI-enabled technologies, any national AI R&D strategy will need to be frequently revisited—in light of regulatory learning and the changing AI landscape in the United States and other major jurisdictions. An adaptable and light-touch regulatory approach is needed to secure America's global economic competitiveness and technological innovation in AI and emerging technologies that depend on artificial intelligence.

Ryan Nabil
Research Fellow
Competitive Enterprise Institute

⁴² For a longer discussion, see Hilary Allen, “Regulatory Sandboxes.” *The George Washington Law Review* 87, no. 3 (May 2019): 579–645. <http://dx.doi.org/10.2139/ssrn.3056993>.

⁴³ Ho et al, *Building a National AI Research Resource*, pp. 28–33.

⁴⁴ *Ibid.*