

September 8, 2022

No. 282

Strategies to Improve the National Artificial Intelligence Research and Development Strategic Plan

By *Ryan Nabil**

Few policy issues are as crucial to the future of U.S. national security and global economic competitiveness as the development of artificial intelligence (AI) and AI-enabled emerging technologies. Artificial intelligence innovation promises to bring considerable consumer benefits in sectors as diverse as finance, legal services, retail and e-commerce, and health care and medical technology. Accordingly, successive U.S. administrations have recognized the need to create a more favorable research environment that enables greater AI innovation in such fields.¹

The White House Office of Science and Technology Policy (OSTP)—in collaboration with the National Science and Technology Council (NSTC) and the National AI Initiative Office—in October 2016 developed the National AI Research and Development (R&D) Strategic Plan, which was subsequently updated in June 2019.² As the U.S. government seeks to recalibrate the country’s AI research strategy, it has an opportunity to strengthen America’s position as a global center of AI innovation.³ To accomplish that goal, several reforms are needed to the U.S. government’s current approach toward artificial intelligence research and development.

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 federal R&D spending and grants to various recipients in different AI subdisciplines.

Third, the national strategic plan would benefit from a fuller understanding of other countries’ AI R&D strategies, which could help the U.S. government allocate federal research grants more effectively. To that end, developing mechanisms to monitor rapidly evolving AI-related laws and regulatory frameworks in foreign jurisdictions could help American lawmakers and regulators capitalize on other countries’ successes and avoid potential regulatory mistakes.

* *Ryan Nabil is a Research Fellow at the Competitive Enterprise Institute.*

Fourth, as proposed in the national AI strategy, developing shared AI data sets for public use can help strengthen AI research and innovation at academic institutions and in the private sector.

Fifth, the National AI R&D Strategic Plan should propose a federal AI regulatory sandbox program—to incentivize the private sector and research institutions to play a more important role in AI research. By allowing companies and academic institutions to test innovative AI systems and AI-enabled technologies for a limited time, a federal sandbox can help promote innovation, enhance regulatory understanding of AI, and help craft innovation-friendly regulatory frameworks and technical standards for artificial intelligence.

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.⁴ 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 developing artificial intelligence technologies and provides several recommendations to enhance public-private collaboration. For instance, it proposes creating joint public-private collaboration, increasing the availability of public data sets, and expanding AI training and fellowship opportunities to meet workforce R&D needs.⁵ 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 below.

Developing Mechanisms to Track and Evaluate Artificial Intelligence R&D Spending. Greater transparency and more precise information about federal AI spending and its impact on innovation within different AI subdisciplines can help policy makers allocate R&D resources more effectively. However, despite the growth in federal spending on AI-related research and development, there appears to be a scarcity of efforts in tracking how this money is spent and how it might impact AI innovation.⁶

Therefore, 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. Collecting and analyzing data to answer some important questions can help policy makers make more informed, evidence-based spending decisions. Have resources allocated to specific AI subdisciplines—such as computer vision and natural language processing—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?

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.⁷

However, tracking spending can nonetheless help compare the effectiveness of similar short- and long-term projects by different agencies, research institutions, and companies. Such data could enable the U.S. government to allocate more resources to promising AI subdisciplines and help improve competition among recipients of federal research grants to submit the most promising proposals.

Lessons from the Successes and Shortcomings of Other Countries' AI Policies. The National AI Strategic R&D Plan would benefit from a closer examination of how other countries allocate research spending, their regulatory approach toward AI, and the extent to which these policies have been successful, or not. 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.

AI policies and developments in other countries also provide opportunities to better understand which R&D and regulatory approaches have been successful elsewhere, and which have not. Policy makers should exercise caution in making such comparisons, as other jurisdictions' regulatory experiences might have limited applicability to the United States. Still, awareness of those broader trends can help the U.S. capitalize on different countries' successes and avoid their regulatory mistakes. To maximize the benefit of this comparative approach, the strategic plan should propose mechanisms to conduct periodic reviews of the global AI research and regulatory landscape and to evaluate international successes and failures.

For example, the National AI R&D Strategic Plan could benefit from a closer examination of European and Chinese approaches toward AI research. While the European Union's overly restrictive approach to AI risks harming innovation in certain sectors, several EU member countries have designed innovative AI R&D strategies at the national level.⁸ 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.⁹ 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.¹⁰

Studying and evaluating these countries' approaches might provide American policy makers 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 launch an AI sandbox and create high-quality public data sets 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.

It is equally important to examine other countries' regulatory mistakes alongside their successes. For instance, as part of their effort to influence global AI norms, the United Kingdom and the European Union have sought to promote algorithmic transparency for AI-enabled technologies.¹¹ However, as discussed later, such proposals have been met with concerns from British and European policy makers and institutions since they could negatively affect technological innovation.¹² Likewise, some American policy makers and experts advocate transparency as a goal in developing AI systems.¹³ 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.¹⁴

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.¹⁵ 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 were to be 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 preventing cyberattacks and detecting diseases accurately are much more important than explaining the underlying algorithms of AI systems.

Therefore, the National AI R&D Strategic Plan should not prioritize transparency over accuracy as a research objective—except in limited, prespecified cases such as AI systems for criminal justice, law enforcement, and human resources.¹⁶

Moreover, mandating high levels of algorithmic transparency from such programs can detract from cutting-edge AI innovation in fields from health care to nuclear science. AI systems typically rely on complex neural networks and algorithms, so their underlying processes are not transparent even to their designers. 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 and that is reported to be substantially better than other comparable systems. Deep Patient, which trained on the medical data of 700,000 patients across several hundred variables, can accurately provide medical diagnostics, but its engineers cannot accurately describe how the program arrives at such a diagnosis.¹⁷

Nevertheless, despite opposition from government officials and data scientists, some policy makers in the United Kingdom and at the European Commission are increasingly advocating algorithmic transparency as a global AI norm.¹⁸ 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.¹⁹

Studying these international regulatory developments and criticisms can help American policy makers avoid potential mistakes, such as 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. policy makers maintain an adaptable, innovation-friendly approach toward artificial intelligence.

Encouraging Academic and Private Sector Innovation by Providing Access to Non-Sensitive Government Databases. 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 large technology companies. Creating innovative AI systems requires high-quality data sets on which AI systems can be trained. The costs associated with creating, cleaning, and preparing such data sets remain too high for many businesses and academic institutions.²⁰ For example, London-based Google subsidiary DeepMind’s AlphaGo software made headlines in March 2016 when it defeated the human champion of the Chinese game Go.²¹ The cost to train data sets for building this program was more than \$25 million in hardware alone.²²

Recognizing this challenge, the National AI R&D Strategic Plan recommended the development of shared data sets that startups, businesses, and research institutions can use to create and train AI programs.²³ Despite this commitment, progress in this area has been slow. That is why the AI strategic plan needs to outline more concrete steps to publish high-quality data sets 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 entail.²⁴

However, different U.S. agencies also have access to a wide range of non-personally identifiable and non-sensitive data sets intended for public use—such as the National Oceanic and Atmospheric Administration’s climate data and the National Aeronautics and Space Administration’s non-confidential space-related data.²⁵ Making such data readily available to the public can allow AI innovation in weather forecasting, transportation, astronomy, and other areas. Therefore, the national 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 deciding which data should be included in public data sets. 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, such as personally identifiable information, which are controlled, unclassified data unavailable to the public; and
- 3) High-impact risk data containing “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 data sets 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 safeguards so that these data sets do not erroneously contain sensitive information, uphold cybersecurity best practices, and are provided in a format suitable for training AI systems.

Creating a Federal Artificial Intelligence Regulatory Sandbox Program. The National AI R&D Strategic Plan should propose the creation of a federal AI sandbox program to develop the regulatory understanding of emerging AI technologies and incentivize private sector innovation.

Given the rapid evolution of AI-enabled technologies, there is a growing need to better understand their ethical, legal, and societal implications.²⁷ As the strategic plan notes, “the unusual complexity and evolving nature” of AI systems mean that “ensuring the safety of AI systems” remains a significant challenge.²⁸ Understanding such systems would constitute an essential first step toward crafting safe, market-friendly regulatory frameworks and technical standards for AI systems and AI-enabled technologies.²⁹

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 They can also help lawmakers and regulators better understand emerging technologies and develop innovation-friendly regulatory frameworks accordingly.

The United Kingdom’s Financial Conduct Authority created the world’s first financial technology (FinTech) sandbox program in 2015. Since then, regulators in Australia, Singapore, Switzerland, and a growing number of jurisdictions have launched similar programs.³¹ In the U.S., the federal Consumer Financial Protection Bureau and at least 11 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 legal services.³³

While FinTech sandbox programs are becoming increasingly common, AI sandboxes remain largely underexplored.³⁴ In its 2021 Artificial Intelligence Act, the European Commission first proposed the creation of coordinated national-level AI sandbox programs to promote AI innovation across Europe.³⁵ In June 2022, Spain became the first EU member to launch a national AI sandbox, and other European countries, including France and the Czech Republic, could soon follow suit.³⁶ The success of such sandboxes will

depend on their regulatory design and the extent to which they prioritize technological innovation and regulatory learning.

In the United States, the updated National AI R&D Strategic Plan could recommend the creation of a federal AI sandbox that accepts participants based on proposed innovation and its potential benefit to developing novel AI-enabled technologies.³⁷

Regulators could also create sandbox programs to target innovation in specific areas—such as human-machine interaction and probabilistic reasoning—that the strategic plan identifies as areas in need of further research.³⁸ For example, current AI systems are ill-equipped to translate heavily accented speech or speech in a noisy environment.³⁹ 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-based cybersecurity and encryption tools can help to encourage market innovations to counter growing cybersecurity challenges.⁴⁰

Additionally, regulators will need to define the type of AI systems and AI-enabled products and services eligible for participating in the sandbox. For example, 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, which is usually one to two years.⁴¹

Another challenge in sandbox design is the fact that the systems, products, and services eligible for the sandbox could fall under the jurisdictions of multiple regulators. Therefore, an AI sandbox might require a well-crafted legal framework that clarifies 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. Unlike in the financial services industry, much of AI innovation relies on academic research.⁴² Typically, FinTech sandbox programs provide participants with exemptive regulatory relief and waived or expedited registration processes.⁴³ However, some research institutions might not seek to commercialize the algorithms and technologies they test in a sandbox, so exemptive regulatory relief might be of limited benefit to them.

Furthermore, for many research institutions and companies, the lack of access to high-quality data sets and cloud-based computing power might pose a greater obstacle than regulatory barriers. 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 U.S. government should focus on enabling a wide range of actors, from tech startups to academic and financial institutions, to play a role in artificial intelligence innovation.

Given rapid technological changes, any national AI R&D approach will need to be frequently revisited—in light of regulatory learning and the evolving AI landscape in the United States and other major jurisdictions. An adaptable and light-touch regulatory approach will help secure America’s global economic competitiveness and technological innovation in artificial intelligence and AI-enabled emerging technologies.

Notes

¹ This policy brief is based on the regulatory comments submitted by the author 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. Ryan Nabil, “Comments in the Matter of the Office of Science and Technology Policy RFI Response: National Artificial Intelligence Research and Development Strategic Plan,” 87 FR 5876, Document Number 2022-02161, https://cei.org/regulatory_comments/cei-comments-on-office-of-science-and-technology-policys-request-for-comments-on-updating-the-national-artificial-intelligence-research-and-development-strategic-plan/.

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, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>.

³ Office of Science and Technology Policy, “Request for Information to the Update of the National Artificial Intelligence Research and Development Strategic Plan,” 87 FR 5876, Document Number 2022-02161, <https://www.federalregister.gov/documents/2022/02/02/2022-02161/request-for-information-to-the-update-of-the-national-artificial-intelligence-research-and->

⁴ NSTC, *The 2019 National AI R&D Strategic Plan*, p. iii. 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>.

⁵ NSTC, *The 2019 National AI R&D Strategic Plan*, p. iii.

⁶ 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).

⁷ 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.

⁸ 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%C2%A0goes-too-far-191733>. Benjamin Mueller, *How Much Will the Artificial Intelligence Act Cost Europe?* (Brussels: Center for Data Innovation, July 2021), <https://www2.datainnovation.org/2021-aia-costs.pdf>.

⁹ Daniel Castro and Joshua New, “Regulatory Comments in Response to the Request for Information on Update to the 2016 National AI R&D Strategic Plan,” Center for Data Innovation, October 26, 2018, p. 3, <https://www2.datainnovation.org/2018-nitrd-ai-r%26d.pdf>.

¹⁰ This point should not be construed as an endorsement of the Chinese government’s current approach toward artificial intelligence. However, American and European lawmakers and regulators could still learn from certain aspects of China’s innovative approach to AI—such as its recent focus on interdisciplinary research projects—when they do not contravene U.S. and European economic and political values. 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, 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/>.

¹¹ 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).

¹² 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>.

¹³ Electronic Privacy Information Center, “State Artificial Intelligence Policy,” 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>.

¹⁴ NSTC, *The 2019 National AI R&D Strategic Plan*, p. 21.

¹⁵ Castro and New, pp. 4–5.

¹⁶ *Ibid.*

¹⁷ *Ibid.*, p. 4. Will Knight, “The Dark Secret at the Heart of AI,” *MIT Technology Review*, April 11, 2017,

<https://www.technologyreview.com/2017/04/11/5113/the-dark-secret-at-the-heart-of-ai/>.

¹⁸ 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.”

¹⁹ 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.

²⁰ 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.

²¹ DeepMind, “AlphaGo,” accessed August 26, 2022, <https://www.deepmind.com/research/highlighted-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.

²² Elizabeth Gibney, “Self-taught AI is best yet at strategy game Go,” *Nature*, October 18, 2017,

<https://www.nature.com/articles/nature.2017.22858>.

²³ NSTC, *The 2019 National AI R&D Strategic Plan*, p. 28.

²⁴ Privacy Act of 1974, 5 U.S.C. § 552a (2012). Ho, King et al., *National AI Research Resource*, p. 35.

²⁵ Ho, King, et al., *National AI Research Resource*, p. 47.

²⁶ *Ibid.* p. 39.

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,” news release, 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*, 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,” 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/?uri=CELEX:52021PC0206>. Luca Bertruzzi, “France seeks more flexibility for regulatory sandboxes in AI regulation,” *Euractiv*, June 16, 2022 (updated June 21, 2022), <https://www.euractiv.com/section/digital/news/france-seeks-more-flexibility-for-regulatory-sandboxes-in-ai-regulation/>. Ministry of Industry and Trade of the Czech Republic, “National Artificial Intelligence Strategy of the Czech Republic,” 2019, pp. 34–35, https://www.mpo.cz/assets/en/guidepost/for-the-media/press-releases/2019/5/NAIS_eng_web.pdf.

³⁶ European Commission, “Launch event for the Spanish Regulatory Sandbox on Artificial Intelligence,” event, June 27, 2022, <https://digital-strategy.ec.europa.eu/en/events/launch-event-spanish-regulatory-sandbox-artificial-intelligence>.

³⁷ For a brief overview of federal and state-level FinTech sandbox programs in the United States, see Ryan Nabil, “How Regulatory Sandbox Programs Can Promote Technological Innovation and Consumer Welfare: Insights from Federal and State Experience,” *OnPoint*, No. 281, Competitive Enterprise Institute, August 17, 2022, <https://cei.org/studies/how-regulatory-sandbox-programs-can-promote-technological-innovation-and-consumer-welfare/>.

³⁸ NSTC, *The 2019 National AI R&D Strategic Plan*, p. 12.

³⁹ *Ibid.*

⁴⁰ 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>.

⁴¹ As the 2019 National AI Strategic Plan notes, 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.

⁴² *Ibid.*, p. iii.

⁴³ For a detailed discussion, see Hilary Allen, “Regulatory Sandboxes,” *George Washington Law Review*, Vol. 87, No. 3 (May 2019), pp. 579–645, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3056993.

⁴⁴ Ho, King, et al, pp. 28–33.