Federal Register Notice 86 FR 46278, https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Stanford University
Human-Centered
Artificial Intelligence

October 1, 2021

Lynne Parker, Ph.D.
Director, National AI Initiative
White House Office of Science and Technology Policy

Erwin Gianchandani, Ph.D.
Senior Advisor, Translation, Innovation, and Partnerships
National Science Foundation

Dear Dr. Parker and Dr. Gianchandani,

We are pleased to submit the attached in response to the Request for Information by the White House Office of Science and Technology and National Science Foundation for creating and implementing a National Artificial Intelligence Research Resource (NAIRR).

Stanford's Institute for Human-Centered Artificial Intelligence (HAI) has long championed the creation of the NAIRR. In 2020, Stanford HAI led efforts with 22 top universities and a bipartisan, bicameral group of lawmakers to pass legislation establishing the NAIRR Task Force. In 2021, we convened a policy practicum course at Stanford on *Creating a National Research Cloud*, which brought together law, business, and engineering students, researchers, and faculty to address some of the core questions for building the NAIRR. To prepare the Report, we interviewed a wide range of stakeholders, engaged in extensive research on existing models, and developed a set of recommendations on both the data and compute dimensions.

To adhere to the 10-page limitation for submissions, please find enclosed the Executive Summary from our Report. Because our full Report, "Building a National AI Research Resource: A Blueprint for the National Research Cloud," spans 100+ pages, we have enclosed a draft as supplemental material. We anticipate that the final Report will be made publicly available the week of October 4.

We hope our extensive work on the design of the NAIRR proves useful to the Task Force and we are happy to engage the Task Force in any way we might be helpful.

Sincerely,

Daniel E. Ho, Ph.D.
William Benjamin Scott and Luna M. Scott Professor of Law
Associate Director, Stanford HAI

Jennifer King, Ph.D.
Privacy and Data Policy Fellow, Stanford HAI

Russell C. Wald
Director of Policy, Stanford HAI

Christopher Wan
JD/MBA Candidate, Stanford University

Executive Summary: Creating a National Research Cloud[1]

Artificial intelligence (AI) appears poised to transform the economy across sectors ranging from healthcare and finance, to retail and education. What some have coined the "Fourth Industrial Revolution"[2] is driven by three key trends: greater availability of data, increases in computing power, and improvements to algorithm design. First, increasingly large amounts of data have fueled the ability for computers to learn, such as by training an algorithmic language model on all of Wikipedia.[3] Second, better computational capacity (often termed "compute") and compute capability have enabled researchers to build models that were unimaginable merely 10 years ago, sometimes spanning billions of parameters (an exponential increase in scope from previous models).[4] Third, basic innovations in algorithms are helping scientists to drive forward AI, such as the reinforcement learning techniques that enabled a computer to defeat the world champion in the board game Go.[5]

Despite these trends, the AI innovation landscape faces serious potential challenges. Historically, partnerships between government(s), universities, and industries have anchored the U.S. innovation ecosystem. The federal government played a critical role in subsidizing basic research, enabling universities to undertake high-risk research that can take decades to commercialize. This approach catalyzed radar technology, the internet, and GPS devices. As the economists Ben Jones and Larry Summers put it, "[e]ven under very conservative assumptions, it is difficult to find an average return below $4 per $1 spent" on innovation, and the social returns might be closer to $20 for every dollar spent. [6] Industry in turn, scales and commercializes applications.

Core challenges to this ecosystem and the future of AI exist. Computing power has become critical for the advancement of AI, but the high cost of compute has placed cutting edge AI research in a position accessible only to key industry players and a handful of elite universities.[7] Access to data—the raw ingredients used to train most AI models—is increasingly

---

[1] This is the Executive Summary to a broader report on the National Research Cloud: DANIEL. E. HO, JENNIFER KING, RUSSELL C. WALD, CHRISTOPHER WAN, BUILDING A NATIONAL AI RESEARCH RESOURCE: A BLUEPRINT FOR THE NATIONAL RESEARCH CLOUD (2021).

[2] KLAUS SCHWAB, THE FOURTH INDUSTRIAL REVOLUTION (2016).

[3] Tae Yano & Moonyoung Kang, *Taking Advantage of Wikipedia in Natural Language Processing*, CARNEGIE MELLON U. (2008), https://www.cs.cmu.edu/~taey/pub/wiki.pdf.

[4] *See, e.g.,* Anthony Alford, *Google Trains Two Billion Parameter AI Vision Model,* INFOQ (June 22, 2021), https://www.infoq.com/news/2021/06/google-vision-transformer/*;* Anthony Alford, *OpenAI Announces GPT-3 AI Language Model with 175 Billion Parameters,* INFOQ (June 2, 2020), https://www.infoq.com/news/2020/06/openai-gpt3-language-model/.

[5] *AlphaGo*, DEEPMIND (2021), https://deepmind.com/research/case-studies/alphago-the-story-so-far/.

[6] Benjamin F. Jones & Lawrence H. Summers, *A Calculation of the Social Returns to Innovation* (Nat'l Bureau of Econ. Research, Working Paper No. 27863, 2020); J.G. Tewksbury, M.S. Crandall & W.E. Crane, *Measuring the Societal Benefits of Innovation*, 209 SCI. MAG. 658-62 (1980); *see also* NATIONAL ACADEMIES OF SCIENCES, ENGINEERING, AND MEDICINE, RETURNS TO FEDERAL INVESTMENTS IN THE INNOVATION SYSTEM (2017)

[7] STUART ZWEBEN & BETSY BIZOT, 2019 TAULBEE SURVEY: TOTAL UNDERGRAD CS ENROLLMENT RISES AGAIN, BUT WITH FEWER NEW MAJORS; DOCTORAL DEGREE PRODUCTION RECOVERS FROM LAST YEAR'S DIP (2019).

limited to the private sector and large platforms[8], since government data sources remain largely inaccessible to the AI research community.[9] As the National Security Commission on AI (NSCAI) has determined, "[t]he consolidation of the AI industry threatens U.S. technological competitiveness."[10] Four interrelated challenges illustrate this finding: First, we are seeing a significant brain drain of researchers departing universities.[11] In 2011, AI Ph.D.s were roughly equally likely to go into industry vs. academia.[12] Ten years later, two- thirds of AI Ph.D.s go to industry, and less than one quarter go into academia.[13] Second, these trends indicate that many university researchers struggle to engage in cutting-edge science, draining the field of the diverse set of research voices that it needs. Third, the fundamental research that would guarantee the United States stays at the helm of AI innovation is being crowded out. By one estimate, 82 percent of algorithms used today originated from federally funded nonprofits and universities, but "U.S. leadership has faded in recent decades."[14] Fourth, government agencies have faced challenges in building compute infrastructure,[15] and there are societal benefits to reducing the cost of core governance functions and improving government's internal capacity to develop, test, and hold AI systems accountable.[16] In short, a growing imbalance in AI innovation tilts towards industry, leaving academic and non-commercial research behind. Given the longstanding role of academic and non-commercial research in innovation, this shift has substantial negative consequences for the American research ecosystem.

Responding to these challenges, Congress enacted the National AI Research Resource Task Force Act as part of the National Defense Authorization Act (NDAA) in January 2021.[17] The Act forms part of the National Artificial Intelligence Initiative, which identifies further steps to increase research investments, set technical standards, and build a stronger AI workforce. The Act created a Task Force—the composition of which was announced on June 10, 2021[18]—to study and plan for the implementation of a "National Artificial Intelligence Research Resource," namely "a system that provides researchers and students across scientific fields and disciplines with access to compute resources, co-located with publicly available, artificial intelligence-ready

---

[8] Jathan Sadowski, *When Data is Capital: Datafication, Accumulation, and Extraction*, 2019 BIG DATA & SOC'Y 1 (2019).
[9] Amy O'Hara & Carla Medalia, *Data Sharing in the Federal Statistical System: Impediments and Possibilities*, 675 ANNALS AM. ACAD. POL. & SOC. SCI. 138, 140-41 (2018).
[10] NAT'L SECURITY COMM'N ON ARTIFICIAL INTELLIGENCE, FINAL REPORT 186 (2021).
[11] STAN. U. INST. FOR HUMAN-CENTERED ARTIFICIAL INTELLIGENCE, 2021 ARTIFICIAL INTELLIGENCE INDEX REPORT 118 (2021).
[12] *Id.*
[13] *Id.*
[14] Neil C. Thompson, Shuning Ge & Yash M. Sherry, *Building the Algorithm Commons: Who Discovered the Algorithms that Underpin Computing in the Modern Enterprise?*, 11 GLOBAL STRATEGY J. 17-33 (2020).
[15] *See, e.g.*, U.S. GOV'T ACCOUNTABILITY OFFICE, FEDERAL AGENCIES NEED TO ADDRESS AGING LEGACY SYSTEMS (2016); U.S. GOV'T ACCOUNTABILITY OFFICE, CLOUD COMPUTING: AGENCIES HAVE INCREASED USAGE AND REALIZED BENEFITS, BUT COST AND SAVINGS DATA NEED TO BE BETTER TRACKED (2019).
[16] DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 6, 71-72 (2020).
[17] William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5106.
[18] *The Biden Administration Launches the National Artificial Intelligence Research Resource Task Force*, THE WHITE HOUSE (June 10, 2021), https://www.whitehouse.gov/ostp/news-updates/2021/06/10/the-biden-administration-launches-the-national-artificial-intelligence-research-resource-task-force/.

government and non-government data sets."[19] This research resource has also been referred to as the National Research Cloud (NRC) and was strongly endorsed by the NSCAI, which wrote that the NRC "will strengthen the foundation of American AI innovation by supporting more equitable growth of the field, expanding AI expertise across the country, and applying AI to a broader range of fields."[20]

While other initiatives have sought to improve access to compute or data in isolation,[21] the NRC will generate distinct positive externalities by integrating compute and data, the two bottlenecks for high-quality AI research. Specifically, the NRC will provide affordable access to high-end computational resources, large-scale government datasets in a secure cloud environment, and the necessary expertise to benefit from this resource through a close partnership between academia, government, and industry. By expanding access to these critical resources in AI research, the NRC will support basic scientific AI research, the democratization of AI innovation, and the promotion of U.S. leadership in AI.

Stanford Law School's policy practicum program convened a multidisciplinary research team of graduate students, staff, and faculty drawn from Stanford's business, law, and engineering schools to study the feasibility of, and considerations for designing the NRC. Over the past six months, this group studied existing models for compute resources and government data, interviewed a wide range of government, computer science, and policy experts, and examined the technical, business, legal, and policy requirements. This report was commissioned by Stanford's Institute for Human-Centered Artificial Intelligence (HAI), which originated the proposal for the NRC in partnership with 21 other research universities.[22]

Throughout our research, we observed three primary themes that cut across all areas of our investigation. We have integrated these themes into each section of our report and drawn on them to explain our findings.

- *Complementarity between compute and data*. As we evaluated the existing computing and data-sharing ecosystems, one of the systemic challenges we observed was a decoupling of compute resources from data infrastructures. High-performance computing can be useless without data; and a major impediment to data sharing, particularly for high-value government data, lies in requirements for a secure, privacy-protecting computing environment.
- *Rebalancing AI research toward long-term, academic, and non-commercial research.* Presently, AI innovation is disproportionately dependent on the private sector. Public

[19] William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5107 (g).
[20] NAT'L SECURITY COMM'N ON ARTIFICIAL INTELLIGENCE, *supra* note 10, at 191.
[21] *See, e.g.*, *Cloudbank*, https://www.cloudbank.org; *Fact Sheet: National Secure Data Service Act Advances Responsible Data Sharing in Government*, DATA COALITION (May 13, 2021), https://www.datacoalition.org/fact-sheet-national-secure-data-service-act-advances-responsible-data-sharing-in-government/.
[22] Steve Lohr, *Universities and Tech Giants Back National Cloud Computing Project*, N.Y. TIMES (June 30, 2020), https://www.nytimes.com/2020/06/30/technology/national-cloud-computing-project.html; John Etchemendy & Fei-Fei Li, *National Research Cloud: Ensuring the Continuation of American Innovation*, STAN. U. INST. FOR HUMAN-CENTERED ARTIFICIAL INTELLIGENCE, (Mar. 28, 2020), https://hai.stanford.edu/news/national-research-cloud-ensuring-continuation-american-innovation.

investment in basic AI infrastructure can both support innovation in the public interest and complement private innovation efforts. The NRC directs more resources toward AI development in the public interest and helps ensure long-term leadership by the United States in the field by supporting the kind of pure, basic research that the private sector cannot undertake alone.

- *Coordinating short-term and long-term approaches to creating the NRC*. Our research considers many near-term pathways for standing up a working version of the NRC by spelling out how to work within existing constraints. We also identify the structural, legal, and policy challenges to be addressed in the long term for executing the full vision of the NRC.

We summarize our main recommendations here.

Compute Model

- The "Make or Buy" Decision. The main policy choice will be whether to build public computing infrastructure or purchase services from existing commercial cloud providers.
    - It is well established that, based solely on hardware costs, it is more cost-effective to own infrastructure when computing demand is close to continuous.[23] The government also has experience building high-performance computing clusters, typically built by contractors and operated by national laboratories.[24] The National Science Foundation (NSF) has also supported many supercomputing initiatives at academic institutions.[25]
    - The main countervailing concerns are that existing commercial cloud providers have software stacks and usability that AI researchers have widely adopted and may consider to be a more user-friendly platform. Commercial cloud providers offer a way to expand capacity expeditiously, although scale and availability will still be constrained by the availability of current graphics processing unit (GPU) computing resources.
    - We hence recommend a dual investment strategy:
        - First, the compute model of the NRC can be quickly launched by subsidizing and negotiating cloud computing for AI researchers with existing vendors, expanding on existing initiatives like the NSF's CloudBank project.[26]
        - Second, the NRC should invest in a pilot for public infrastructure to assess the ability to provide similar resources in the long run. Such publicly owned infrastructure would still be built under contract or grant, but could

[23] Jennifer Villa & Dave Troiano, *Choosing Your Deep Learning Infrastructure; The Cloud vs. On-Prem Debate*, DETERMINED AI (July 30, 2020), https://determined.ai/blog/cloud-v-onprem/; *Is HPC Going to Cost Me a Fortune?*, INSIDEHPC (last visited July 23, 2021), https://insidehpc.com/hpc-basic-training/is-hpc-going-to-cost-me-a-fortune/.
[24] *See, e.g., US Plans $1.8 Billion Spend on DOE Exascale Supercomputing*, HPCWIRE (Apr. 11, 2018), https://www.hpcwire.com/2018/04/11/us-plans-1-8-billion-spend-on-doe-exascale-supercomputing/; *Federal Government*, ADVANCED HPC (last visited July 23, 2021), https://www.advancedhpc.com/pages/federal-government; *United States Continues to Lead World In Supercomputing*, U.S. DEP'T. ENERGY (Nov. 18, 2019), https://www.energy.gov/articles/united-states-continues-lead-world-supercomputing.
[25] *See NSF Funds Five New XSEDE-Allocated Systems,* NAT'L SCI. FOUND. (Aug. 10, 2020), https://www.xsede.org/-/nsf-funds-five-new-xsede-allocated-systems.
[26] *Cloudbank*, *supra* note 21.

be operated much like national laboratories (e.g., Sandia National Laboratories, Oak Ridge National Laboratory), that own sophisticated supercomputing facilities or academic supercomputing facilities.

- Researcher Eligibility. While some have argued the NRC should be open for commercial access, for the purposes of this report, we adhered to the spirit of the legislation forming the NAIRR Task Force and only reviewed the use of an NRC for academic and non-profit AI research. We recommend that the NRC eligibility start with academics who hold "Principal Investigator" (PI) status (i.e., most faculty) at U.S. colleges and universities, as well as to "Affiliated Government Agencies" willing to contribute previously unreleased, high-value datasets to the NRC in return for subsidized compute resources. PI status should be interpreted expansively to encompass all fields of AI application. Students working with PIs should presumptively gain access to the NRC. Scaling the NRC to meet the demand of all students in the United States may be challenging, but we also recommend the creation of educational programs as part of the new resource to help train the next generation of AI researchers.
- Mechanism. In order to keep the award processing costs down, we recommend a base-level of compute access to meet the majority of researcher computing needs. Base-level access avoids high overhead for grant administration and may meet the compute demands for the supermajority of researchers. For researchers with exceptional needs, we recommend a streamlined grant process for additional compute access.

Data Access Model

- Focus on Government Data. We focus our recommendations for data provision/access to government data because: (1) there are already a wide range of platforms for sharing private data,[27] and (2) distribution by the NRC of private datasets would raise a tangle of thorny IP issues. We recommend that researchers be allowed to compute on any datasets they themselves contribute, provided they certify they have the rights to that data, and the use of such data is for academic research purposes.
- Tiered Access. We recommend a tiered access model: by default, researchers will gain access to government data that is already public; researchers can then apply through a streamlined process to gain access at higher security levels on a project-specific basis. It will be critical for the NRC to ultimately displace the current fragmented, agency-by-agency relational approach. By providing secure virtual environments and harmonizing security standards (e.g., Federal Risk and Authorization Management Program (FedRAMP)[28]), the NRC can collaborate with proposals for a National Secure Data Service[29] to provide a model for accelerating AI research, while protecting data privacy and prioritizing data security.
- Agency Incentives. To incentivize federal agencies to share data with the NRC and improve the state of public sector technology, we recommend the NRC permit federal

[27] *See, e.g., National Data Service*, http://www.nationaldataservice.org; *The Open Science Data Cloud*, https://www.opensciencedatacloud.org; *Harvard Dataverse*, https://dataverse.harvard.edu; *FigShare*, https://figshare.com.
[28] *FedRAMP*, https://www.fedramp.gov.
[29] *See Fact Sheet: National Secure Data Service Act Advances Responsible Data Sharing in Government,* DATA COALITION (May 13, 2021), https://www.datacoalition.org/fact-sheet-national-secure-data-service-act-advances-responsible-data-sharing-in-government/.

agency staff to use the NRC's compute resources. In keeping with the practices of existing data-sharing programs, such as the Coleridge Initiative,[30] we also recommend that the NRC provide training and support to work with agencies to modernize and harmonize their data standards.

- Strategic Investment for Data Sources. In the short term, we recommend that the NRC focus its efforts on making available non-sensitive, low- to moderate-risk government datasets, rather than sensitive government data (e.g., data about individuals) or data from the private sector, due to data privacy and intellectual property concerns. Researchers can still use NRC compute resources on private data, but should rely on existing mechanisms to acquire data for their own private buckets on the NRC. For example, images taken from Earth observation satellites, such as Landsat imagery, provide a promising low-risk, high-reward government dataset, as making such satellite imagery freely available to researchers has generated an estimated $3-4 billion in annual economic benefits, particularly when combined with high-performance computing.[31] Agencies such as the National Oceanic and Atmospheric Administration, the U.S. Geological Survey, the Census Bureau, the Administrative Office of the U.S. Courts, and the Bureau of Labor Statistics, for instance, also have rich datasets that can more readily be deployed. In the long run, access to high-risk datasets, such as those owned by the Internal Revenue Service (IRS) and the Veterans Administration (VA), will depend on the tiered access model.

Organizational Form

Where to institutionally locate the NRC poses a tradeoff between ease of coordination to obtain compute and ease of data access. For instance, locating the NRC within a single agency would make coordination with compute providers easier, but would make data access across agencies more difficult, absent further statutory authority. Many efforts to make data access to government data easier, most notably the Foundations for Evidence-Based Policymaking Act of 2018, have proven to be among the most daunting challenges of government modernization.[32] Building on those insights, we ultimately recommend that the NRC be instituted as a Federally Funded Research and Development Center (FFRDC) in the short run, and a public-private partnership (PPP) in the long run.

- FFRDC. FFRDCs at Affiliated Government Agencies would reduce the significant costs of securing data from those host agencies. This approach will also cohere with the greater reliance on commercial cloud credits in the short run, making compute and data

[30] *See Administrative Data Research Facility*, COLERIDGE INITIATIVE, https://coleridgeinitiative.org/adrf/ (last visited July 26, 2021).
[31] *See Landsat Data Access*, U.S. GEOLOGICAL SURVEY, https://www.usgs.gov/core-science-systems/nli/landsat/landsat-data-access (last visited July 23, 2021); FED. GEOGRAPHIC DATA COMM., THE VALUE PROPOSITION FOR LANDSAT APPLICATIONS (2014); CRISTA L. STRAUB, STEPHEN R. KOONTZ & JOHN B. LOOMIS, ECONOMIC VALUATION OF LANDSAT IMAGERY (2019).
[32] *See* BIPARTISAN POL'Y CTR., BARRIERS TO USING GOVERNMENT DATA: EXTENDED ANALYSIS OF THE U.S. COMMISSION ON EVIDENCE-BASED POLICYMAKING'S SURVEY OF FEDERAL AGENCIES AND OFFICES 18-20 (2018); *see also* U.S. DEP'T OF HEALTH & HUMAN SERVICES, THE STATE OF DATA SHARING AT THE U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES 4 (2018) (describing how data at the agency is "largely kept in silos with a lack of organizational awareness of what data are collected across the Department and how to request access.").

coordination less central. In the long run, however, streamlined coordination between data and compute may be more difficult with FFRDCs hosted at specific agencies when (1) the NRC moves away from commercial cloud credits and towards its own high-performance computing cluster, and (2) a greater number of inter-agency datasets become available.

- PPP. In the long run, we recommend the creation of a PPP model, governed by officers from Affiliated Government Agencies, academic researchers, and representatives from the technology sector, which can house both compute and data resources.

Additional Considerations

- Data Privacy. As an initial matter, an NRC where sensitive or individually identifiable administrative data from multiple agencies are used to build and train AI models will face challenges from the Privacy Act of 1974.[33] The Act is intended to put a check on interagency data sharing and disclosure of sensitive data without consent.
    - In order to avoid conflicts with non-consensual interagency data sharing, we recommend that the NRC should not be instituted as its own federal agency, nor should federal agency staff be allowed access to interagency data.
    - To avoid conflicts with the Act's "no disclosure without consent" requirement, any data released to the NRC must not be individually identifiable. Despite these constraints, the majority of AI research will likely fall under the Act's statistical research exception, contingent on proposals aligning with an agency's core purpose.
    - Given concerns about the potential privacy risks, federal agencies may desire to share data, contingent on the use of technical privacy measures (e.g., differential privacy). While useful in many instances, technical approaches are no panacea and should not substitute for data access policies.
    - The NRC should explore the design of virtual "data safe rooms" that enable researchers to access data in a secure, monitored, and cloud-based environment.
    - Additional legislative interventions could also facilitate data sharing with the NRC (e.g., requiring IT modernization to include data sharing plans with the NRC).
- Ethics. Rapid innovation in AI research raises a host of potential ethical challenges. Given the scope of the NRC, it will be infeasible to review every single research proposal for potential ethical violations, particularly since ethical standards are still in flux. The NRC should adopt a twofold approach.
    - First, for default PI access to base-level data and compute, the NRC should establish an ex-post review process for allegations of ethical research violations. Access may be revoked when research is shown to manifestly and seriously violate ethical standards. We emphasize that the high standard for a violation should be informed by the academic speech implications and potential political consequences of government involvement in administering the NRC and determining academic research directions.
    - Second, for applications requesting access to restricted datasets or resources beyond default compute, which will necessarily undergo some review, researchers

[33] Privacy Act, 5 U.S.C. § 552a (1974).

should be required to provide an ethics impact statement. One of the advantages of beginning with PIs is that university faculty are accountable under existing IRBs for human subjects research, as well as to the tenets of peer review.
    - We urge non-NRC parties (e.g., universities) to explore a range of measures to address ethical concerns in AI compute (e.g., an ethics review process[34] or embedding ethicists in projects[35]).
- Security. We recommend that the NRC take the lead in setting security classifications and protocols, in part to counteract a balkanized security system across federal agencies that would stymie the ability to host datasets. The NRC should use dedicated security staff to work with Affiliated Government Agencies and university representatives to harmonize and modernize agency security standards.
- Intellectual Property (IP). While the evidence on optimal IP incentives for innovation is mixed, we recommend that the NRC adopt the same approach to allocating patent rights, copyrights, and data rights to NRC users that apply to federal funding agreements. The NRC should additionally consider conditions for requiring NRC researchers to disclose or share their research outputs under an open-access license.
- Human Resources. Given its ambition, significant human resources – from systems engineers to data officers, and from grants administrators to privacy, ethics, and cybersecurity staff – will be necessary to make the NRC a success.

[34] Michael S. Bernstein et al., *ESR: Ethics and Society Review of Artificial Intelligence Research*, CORNELL. U. (July 9, 2021), https://arxiv.org/pdf/2106.11521.pdf.
[35] Courtenay R. Bruce et al., *An Embedded Model for Ethics Consultation: Characteristics, Outcomes, and Challenges*, 5 AJOB EMPIRICAL BIOETHICS 8 (2014).