

October 1, 2021

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

Dear Members of the National AI Research Resource Task Force,

The AI Now Institute of New York University and Data & Society Research Institute are pleased to submit a response to the Request for Information (RFI) published by the Office of Science and Technology Policy and the National Science Foundation (NSF) to inform the work of the National Artificial Intelligence Research Resource (NAIRR) Task Force.

Our organizations are interdisciplinary research institutes working to help ensure that artificial intelligence (AI) systems are accountable to the communities and contexts in which they are applied, and to produce empirical research that challenges the power asymmetries created and amplified by technology in society. We have worked extensively within academic institutions, civil society and advocacy communities, and in solidarity with marginalized communities and workers directly affected by algorithmic harms.

We are pleased to see the government's commitment to supporting research on AI, and for the opportunity to contribute to the development of the Task Force's implementation blueprint. We are filing this comment to express our concern about the stated aims for the NAIRR and to recommend alternative policy strategies for supporting research into AI and expanding access to data, resources, educational opportunities, and meaningful mechanisms that can ensure democratic oversight of AI and related technologies. We argue that these changes are needed to enable high-quality, interdisciplinary work that goes well beyond a narrow technocratic frame.

Throughout our comment, our recommendations encourage the Task Force to fundamentally reconsider whether the investment in shared computing and data infrastructure is consistent with both the broader aims of the project to “democratize” AI research and the Biden Administration’s explicit commitment to challenging the concentrated power of the tech industry.

In encouraging the Task Force to reconsider, we emphasize that the only plausible short- to mid-term scenario is that the infrastructure required for the NAIRR would be licensed from the same large tech companies responsible for concentrating tech power. This means that while the NAIRR proposal claims to “democratize” access to these resources as a way of contending with the concentrated power of the companies that control AI infrastructures, it would in reality almost certainly work to expand and entrench the power and control these companies have over these

infrastructures.

The NAIRR proposal will make it much harder to check the power of these companies through regulation and public pressure. Furthermore, this proposal assumes a narrow understanding of the AI research field and what kinds of resources are necessary to do the work that's most needed to not only deepen our understanding of AI, but to prevent and mitigate harms that AI and the industry concentration behind it are already causing in our society.

These are challenging issues for which there are no current best practices. Our comment urges the NAIRR Task Force to recommend that Congress consider alternative spending priorities and deeper engagement on the harms unleashed by uncritical investment in AI.

In this comment, we argue:

1. [The NAIRR will entrench, rather than challenge, corporate control over the AI field, contrary to the Biden Administration's bold stance against the power of large tech companies in society.](#)
2. [The NAIRR Task Force must reckon with mounting evidence of the harmful impacts of large-scale AI systems, including discriminatory consequences for marginalized groups and long-term climate impact.](#)
3. [The NAIRR Task Force must expand their understanding of what disciplines constitute "AI research" and redirect NSF resources and programs toward constructing mechanisms for meaningful democratic control of AI and related technologies.](#)
4. [The NAIRR and related proposals raise serious ethical and data privacy challenges, particularly with the use of government data. Given the lack of demonstrated best practices and global policy precedents for data privacy, the NAIRR Task Force should recommend pausing the NAIRR until these challenges are resolved.](#)
5. [The NAIRR, as presently conceived, bolsters misleading and dangerous "tech cold war" narratives, which reflect the self-interest of Big Tech and the defense contracting industry, without being backed by robust evidence.](#)

Background and context: Understanding the "AI boom" (Question 1, Items A and H)

To understand the complex interplay between the issues the NAIRR proposal raises, it's helpful to recall that the current turn to AI is primarily a product of significantly concentrated corporate resources—namely vast computation, massive data, and the capital required to attract and retain

scarce AI talent.¹ **The so-called “advances” in AI that have been celebrated since the early 2010s were not due to breakthroughs in AI research and innovation. They were predicated on newly available access to powerful computation and to massive amounts of web data.** Then, as now, these are resources that a handful of powerful tech companies have in large supply thanks to ad tech-driven surveillance business models, and that few others can avail themselves of without going through these companies.²

The past decade’s Big Tech-led turn to AI profoundly shaped academic computer science disciplines as well. It served to redirect computer science research toward AI-related questions and approaches favored by these companies. In particular, the influx of money and attention produced a turn toward resource-intensive research and development. Work that could avail itself of expensive and scarce industry computing and data was heralded as “cutting edge.” This created an uneasy and conflicted environment for university AI research, in which the dependence on large tech company funding, infrastructure, and data was recognized by practitioners, but not often openly acknowledged.

The NAIRR proposal’s acknowledgement that all of America’s AI researchers cannot “fully participate in”³ this field is a step forward, allowing us to recognize that concentrated corporate power and resources are constitutive of the current wave of AI development. It also presents us with a set of thorny questions, at the center of which is the question of how we reduce the power and control of the handful of companies currently dominating AI and AI research, and how we ensure that determinations about whether—if at all—AI is developed and deployed are subject to more democratic deliberation. While the NAIRR proposal doesn’t currently address these issues, and indeed threatens to exacerbate them, our hope is that the Task Force can redirect attention and resources to this pressing task.

I. The NAIRR will entrench, rather than counter, corporate control over the AI field, contrary to the Biden Administration’s bold stance against the power of large tech companies in society. (*Question 1, Topics H and D; Questions 5 and 6*)

Although the NAIRR is currently envisioned as “a shared computing and data infrastructure that would provide AI researchers and students across scientific fields with access to a holistic advanced computing ecosystem,” it remains unclear how the NAIRR can deliver on the mandate to build an implementation plan in a way that does not further entrench corporate influence and control over the

¹ While the field of Artificial Intelligence is ostensibly oriented around making machines intelligent, in practice, most AI systems rely on big data - the collection and processing of massive datasets, identifying patterns and probabilities within them and codifying them into a predictive mathematical model. See Meredith Broussard (2018) *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge: MIT Press.

² Meredith Whittaker, “The Steep Cost of Capture,” ACM Interactions, Vol. XXVIII.6 Nov-Dec 2021. Forthcoming.

³ <https://www.federalregister.gov/documents/2021/07/23/2021-15660/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>

AI research field.⁴ What is being proposed is an extension of industry-dependent resources, not the construction of resources that would challenge or reduce the centralized power of the large tech players.

From this perspective, it makes sense that the National Security Commission on Artificial Intelligence (NSCAI), a very conflicted body that is helmed by former Google CEO Eric Schmidt, and populated by executives from Amazon, Microsoft, Oracle, and other large tech companies, would prioritize support for such an endeavor.

There is no scenario in the short or mid-term future where large scale computational resources adequate to the task of expanding access to bigger-is-better AI research resources could be created and maintained by institutions meaningfully separate from the large tech platform companies. These companies provide more than raw computing power: the computational environments they own and license provide the tools and research environments that define how AI research gets done. There is no plausible path forward in which such a resource would not be dependent on existing tech industry platforms, tools and resources.

While this isn't stated explicitly, it is tacitly acknowledged in the repeated calls to constitute the NAIRR via "public-private partnerships," a phrase threaded throughout the Stanford Institute for Human-Centered Artificial Intelligence (HAI) and NSCAI's NAIRR advocacy.⁵ NSCAI's recommendations acknowledge that "this infrastructure would leverage public-private partnerships and build on existing government efforts, avoiding high start-up costs of a government-run data center."⁶ Speaking with *Science*, HAI Director John Etchemendy makes the implicit explicit: "The commercial cloud providers are doing the innovation, and they invest massive amounts of money to keep it up-to-date," he says. "It would be a huge mistake to build a facility like a supercomputer center because it would be obsolete within a few years."⁷

These statements make clear to us that the NAIRR proposal is a tacit endorsement of a massive investment in large tech companies in service of expanding their proprietary infrastructure. This arrives at a time when the concentrated power and influence these companies exert is increasingly under scrutiny, including by the Biden Administration itself, which has taken a clear stance that a small number of dominant platforms are using their power to extract monopoly profits.⁸ An effort that aims to "democratize" AI research by investing money in companies that

⁴<https://www.federalregister.gov/documents/2021/07/23/2021-15660/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>

⁵<https://hai.stanford.edu/national-research-cloud-joint-letter>

⁶<https://www.nscai.gov/2020/04/01/nscai-submits-first-quarter-recommendations-to-congress-2/>

⁷<https://www.science.org/news/2021/01/us-law-sets-stage-boost-artificial-intelligence-research>

⁸<https://www.whitehouse.gov/briefing-room/presidential-actions/2021/07/09/executive-order-on-promoting-competition-in-the-american-economy/>

dominate their market will only further entrench these firms' power and reach. This will make it harder to check the power of these companies through regulation and public pressure.

We therefore advise that the NAIRR's forthcoming roadmap and implementation plans recommend against pursuing shared research infrastructure, and instead explore alternative ideas for expanding research into AI, increasing government investment in critical work on AI, and meaningfully democratizing decision making on the development and deployment of AI and related technologies.

II. The NAIRR must reckon with mounting evidence of the harmful impacts of large-scale AI systems, including discriminatory consequences for marginalized groups and long-term climate impact of this scale of computing. (*Question 1, Topics E and G; Question 3*)

The NAIRR Task Force must engage more deeply with the body of critical research, press coverage and investigative reporting, and public discussion that has presented significant evidence of AI's harms, and raised fundamental questions about the ability of AI systems to operate safely and transparently in sensitive social domains. Painting a narrow definition of technical AI research and development as imperative to national success and wellbeing ignores this robust public debate and accompanying evidence, and forecloses discussion about whether society wants these systems in the first place.

While proponents of the NAIRR and the expanding use of AI more generally often point to the potential for this technology to stimulate economic growth, many people—particularly marginalized communities—are already subject to the worst excesses, mistakes, and harms perpetuated by the oppressive and extractive use of powerful algorithmic technology.⁹ Focusing on propelling increasingly unequal economic growth while these harms continue relatively unchecked, raises critical questions about whether certain communities are considered expendable and what kinds of harms are allowable in our society in the name of economic growth.

In a joint letter addressed to OSTP in July 2021, a coalition of civil rights and technology organizations called on the department to center civil rights concerns in AI and technology policy, emphasizing the need to fully incorporate the Biden Administration's Executive Order on Racial Equity into its AI policy priorities. This letter drew on years of evidence about the harms that AI is already causing, including perpetuating housing, financial services, and hiring discrimination.¹⁰ This

⁹ <https://nyupress.org/9781479837243/algorithms-of-oppression/>;
[https://www.wiley.com/en-us/Race+After+Technology:+Abolitionist+Tools+for+the+New+Jim+Code-p-9781509526437#:~:text=Presenting%20the%20concept%20of%20the%20ultimately%20doing%20quite%20the%20opposite.](https://www.wiley.com/en-us/Race+After+Technology:+Abolitionist+Tools+for+the+New+Jim+Code-p-9781509526437#:~:text=Presenting%20the%20concept%20of%20the%20ultimately%20doing%20quite%20the%20opposite.;);
<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

¹⁰ <https://www.aclu.org/news/privacy-technology/how-artificial-intelligence-can-deepen-racial-and-economic-inequities>

letter also calls attention to the lack of civil rights representation on the NAIRR Task Force, despite the enormous implications of the Task Force’s work on civil rights and civil liberties.

Furthermore, many subcategories of AI, especially those that emphasize large-scale data and computing (often called Large Language Models or LLMs), are uniquely prone to perpetuating social harms and entrenching biases.¹¹ These large-scale models, which are trained on troves of internet data from sources like Reddit, exhibit persistent discriminatory outputs.¹² Furthermore, large-scale AI models are built on mass surveillance, which disproportionately impacts marginalized communities,¹³ without implementing meaningful mechanisms for accountability or consent of the public.¹⁴ Their carbon cost is substantial: the amount of processing required to train AI models is both financially and environmentally resource-intensive, and these costs are only likely to expand given industry standards that tie performance metrics to the size of the dataset used to train the model.¹⁵ As a whole, the tech industry is responsible for a global carbon footprint comparable to the aviation industry, and data centers make up 45% of this footprint.¹⁶

Rigorous empirical work to understand ways AI systems and the concentrated power of the companies behind them exacerbate societal harms is happening in a wide range of disciplines and contexts. While this work has helped inform the public and raise urgent questions, it is rarely included in definitions of fundamental AI research and development (R&D), despite how critical it is to our understanding of the feasibility of ideas like the NAIRR and ultimately the tradeoffs of greater adoption of AI in all sectors of our economy. More meaningful public deliberation around these tradeoffs is important for the NAIRR Task Force to pursue, as is additional funding and resources for AI researchers and communities exploring and exposing these harms.

III. The NAIRR Task Force must expand their understanding of what disciplines constitute “AI research” and redirect NSF resources and programs toward constructing mechanisms for meaningful democratic control of AI and related technologies.
(Question 1, Topic I; Questions 3, 4 and 6)

The proposal offers a narrow interpretation of “AI research” by emphasizing infrastructure required for AI research that is reliant on large datasets and massive computational power. In this vision¹⁷, all

¹¹ <https://dl.acm.org/doi/10.1145/3442188.3445922>

¹² <https://arxiv.org/pdf/2101.05783v1.pdf>

¹³ <https://www.law.georgetown.edu/privacy-technology-center/events/color-of-surveillance-2019/>

¹⁴ <https://dl.acm.org/doi/10.1145/3442188.3445922>

¹⁵ <https://arxiv.org/pdf/1907.10597.pdf>

¹⁶ <https://medium.com/@AINowInstitute/ai-and-climate-change-how-theyre-connected-and-what-we-can-do-about-it-6aa8d0f5b32c>

¹⁷ Proponents of the development of the National Research Cloud (NRC) such as Stanford’s HAI also put forward a similarly narrow interpretation of “AI research”.

AI research is “technical,” and research at the cutting edge is that which is most dependent on large datasets and massive computational power.

Although the NAIRR’s vision includes “AI researchers and students across scientific fields,” it’s important to be explicit about the wide range of people, organizations, and disciplines that also contribute to AI R&D. Major contributions to this field are coming from scholars in disciplines like law,¹⁸ anthropology,¹⁹ and history,²⁰ and also from community organizations²¹ gathering qualitative evidence and elevating the lived experience of the people who are surveilled, assessed, and otherwise subject to AI’s determinations and predictions. These disciplines and approaches are most capable of analyzing and addressing the structural issues with AI and related technologies.

Using this more accurate and expansive understanding of the academic disciplines and organizations leading research on AI, the Task Force must consider a wider range of needs than those required by scholars in fields such as machine learning and neuroscience. The Task Force must also deepen its understanding of what it means to “democratize” AI by centering the perspectives and needs of people and organizations whose work concerns AI but does not rely on computing or quantitative data infrastructure.

Despite the NAIRR proposal’s efforts to “democratize access to the cyberinfrastructure that fuels AI research and development,” the infrastructure that would constitute a NAIRR system would almost certainly to be licensed from the same companies responsible for concentrating AI power.

This raises a critically important question: where else could hundreds of millions of dollars be better spent, and how could such funding help us meaningfully ensure democratic control and deliberation over AI and the tech companies responsible?

The NSF has already laid important groundwork to invest in the creation of new research institutes, investing a combined \$220 million in 40 states and in the District of Columbia.²² The NSF can continue to expand this program by substantively funding under-resourced research domains, taking leadership from those most harmed by inequitable uses of AI, and investing in meaningful public control over these powerful technologies. **Rather than routing millions of dollars in taxpayer money to private industry under the guise of public-private partnership, this funding could be**

¹⁸ <https://datasociety.net/library/poverty-lawgorithms/>; <https://ainowinstitute.org/litigatingalgorithms-2019-us.pdf>; <https://www.nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice/>

¹⁹ <https://datasociety.net/library/repairing-innovation/>; <https://global.oup.com/academic/product/predict-and-surveil-9780190684099?cc=us&lang=en&>

²⁰ <https://datasociety.net/library/house-arrest/>; Joy Lisi Rankin (2018) *A People's History of Computing in the United States*, Harvard University Press; Mar Hicks (2017)

²¹ <https://movementalliance.org/about/>, <https://d4bl.org/>, <https://detroitcommunitytech.org/>, <https://www.drivers-united.org/>; <https://mijente.net/2018/10/whos-behind-ice-the-tech-companies-fueling-deportations/>; <http://www.aewoodcenter.org/>; <https://athenaforall.org/>; <https://www.fightforthefuture.org/>

²² https://www.nsf.gov/news/news_summ.jsp?cntn_id=303176

used to create and fund scholarships for underrepresented students and sociotechnical research disciplines, to invest in fellowships that place socio-technical scholars in federal agencies, and to create public engagement and consultation fora that allow communities most harmed by these systems to have a say in their design and deployment. The NSF must also work to preserve the independence of research undertaken in these institutes by eliminating the model in which these institutes are jointly funded by companies.

IV. The NAIRR and related proposals raise serious ethical and data privacy challenges, particularly with the use of government data. Given the lack of demonstrated best practices and global policy precedent, the NAIRR should pause until these issues are resolved and prioritize these issues for further research. (Question 1, Topics E and G; Question 3)

The creation of a National Research Cloud presents significant challenges to ensuring the ethical use of data. As over a decade of work on developing ethical guidelines for the use of open data indicates, these are hard problems. But given the likelihood that the NAIRR will provide access to government data sets, it is critical that there are stringent mechanisms in place to ensure that uses of this data are accountable and ethical, *before* the NAIRR is implemented. NAIRR documentation thus far does not indicate sufficient resources for or attention to ethical and privacy considerations.

There are several potential issues that must be addressed as a baseline by such a proposal if it proceeds with a focus on technical infrastructure. First, there should be some process for evaluating appropriate uses of the NAIRR resources and their potential downstream impact. Already, leading AI research conferences like NeurIPS require researchers to, at a minimum, address the social impact of their work prior to publishing their research findings.²³ We would suggest such a statement be treated as the floor rather than the ceiling for ethics evaluation.

There would also need to be an appropriate method for evaluating proposals. It is important that the actors involved in evaluation have expertise, independence, authority, and most importantly, adequate resources to conduct review of proposals for access, as well as monitor their use once granted. Though the NAIRR provides a development and research environment, any ethical guidelines implemented for access would do little to inhibit downstream harms in deployment, and AI models trained using NAIRR resources could be redeployed and commercialized in applications that depart from what was originally intended. This is a problem that is exacerbated by corporate secrecy, which allows companies to re-deploy academic models without transparency or oversight.

The NAIRR proposal also emphasizes that beneficiaries of these grants could receive access to “high quality government datasets” alongside access to compute. Increasing access to government data,

²³ <https://neuripsconf.medium.com/getting-started-with-neurips-2020-e350f9b39c28>

including anonymized or non-personal datasets, raises persistent data privacy and security concerns, alongside concerns about potential downstream racial and gender bias and inaccuracy.²⁴ While there is active and evolving research around differential privacy and other means to ensure privacy-preserving data disclosures,²⁵ there is a significant gap between theory and practice. Globally, too, this is an area of active policy making²⁶ as evidenced by the European Union’s Data Governance Act (DGA) of 2020 that puts in place foundational norms around tiered access to government data.²⁷ While the DGA stipulates that government agencies can share datasets with private actors “within a secure processing environment” that remains in the control of the public sector, the specifics of how this proposal would be operationalized is yet to be determined, and there are no working precedents to draw from.²⁸

These are all challenging issues for which there are no current best practices nor global policy precedent, and that require deep study and consideration. On this basis, the NAIRR should not proceed as it is currently envisioned. If it does proceed, it must have a clear and evidence-based demonstration that these issues can be addressed.

V. The NAIRR, as presently conceived, bolsters misleading and dangerous “tech cold war” narratives, which reflect the self-interest of Big Tech and the defense contracting industry, without being backed by robust evidence. (Question 1, Topics A and E; Question 6)

Proponents of the NAIRR have lauded the project as being critical to ‘winning’ a global technological race.²⁹ Evoking calls to nationalism,³⁰ these narratives legitimize continued investment in building out the most computationally intensive forms of AI research as a key element of competing with foreign adversaries, particularly China. **Tech company CEOs and former national security officials have been some of the most vocal in endorsing this self-serving narrative, which troublingly echoes cold war framings that served in the past³¹ to accelerate government investment in weapons-related computing.**³² As we have argued, however, this “tech cold war” narrative propels a race to the bottom, where any calls for democratic restraint and regulation of the largest American AI companies is seen as hurting national interest.³³ Indeed, we increasingly see

²⁴ <https://www.cs.princeton.edu/~arvindn/publications/precautionary.pdf>

²⁵ https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3329330

²⁶ https://law.yale.edu/sites/default/files/area/center/china/document/shifting_narratives.pdf

²⁷ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020PC0767>

²⁸ <https://digital-strategy.ec.europa.eu/en/policies/data-governance-act>

²⁹ <https://www.cnas.org/publications/commentary/a-national-cloud-for-all>

³⁰ <https://foreignpolicy.com/2020/08/27/china-tech-facebook-google/>

³¹ <https://online.ucpress.edu/phr/article-abstract/88/4/619/80374/Artificial-Intelligence-and-Japan-s-Fifth>

³² Seymour Melman, *Pentagon Capitalism: The Political Economy of War*, New York, McGraw-Hill, 1970; Paul Edwards, “The Closed World: Computers and the Politics of Discourse in Cold War America” The MIT Press, 1997

³³ <https://medium.com/@AINowInstitute/china-in-global-tech-discourse-2524017ca856>

China-centered arguments against tech accountability and antitrust³⁴ that frame these regulatory interventions as a barrier to national progress in this so-called “AI arms race.”³⁵

The NAIRR, as it is presently conceived, bolsters this misleading and dangerous narrative by legitimizing the consolidation of the AI industry as a foregone conclusion. It not only affirms continued investment in building out computationally intensive AI as a given for national progress, but also further entrenches dependencies on existing AI infrastructure companies like Amazon, Microsoft, Google, and IBM, who unsurprisingly have been among the most vocal proponents. A senior Amazon Web Services executive for example applauded the NAIRR proposal while emphasizing that the AWS Cloud would “increase efficiency while providing unrivaled scalability” to the project. This approach contradicts the Biden Administration’s position that it is United States policy to not tolerate domestic monopolization as the answer to the rising power of foreign monopolies.³⁶

In conclusion, we thank you for your willingness to grapple with the problem of concentrated corporate control of the resources required to produce AI and similar technologies. In light of the depth and complexity of the questions raised by the proposal questions, we encourage the Task Force to fundamentally reconsider whether the NAIRR is consistent with both the broader aims of the project to “democratize” AI research and the Biden Administration’s explicit commitment to challenging the concentrated power of the tech industry. We look forward to supporting the Task Force in developing alternative proposals.

Sincerely,

Amba Kak, Director of Global Policy, AI Now Institute
Brittany Smith, Policy Director, Data & Society
Dr. Sarah Myers West, Post-doctoral Researcher, AI Now Institute
Meredith Whittaker, Faculty Director, AI Now Institute

³⁴https://www.documentcloud.org/documents/21062393-national-security-letter-on-antitrust?responsive=1&title=1&utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axioslogin&stream=top

³⁵<https://www.theverge.com/2019/10/17/20919464/mark-zuckerberg-facebook-china-free-speech-georgetown-tiktok-byte-dance>; <https://techcrunch.com/2019/07/17/facebook-or-china/>

³⁶<https://www.whitehouse.gov/briefing-room/presidential-actions/2021/07/09/executive-order-on-promoting-competition-in-the-american-economy/>