

Source: Michael Dzedzic, Unsplash

By Fernando A. Delgado and Karen Levy

The benefits arising from artificial intelligence (AI) innovations are currently concentrated among only a select few firms,^[1] and the larger ecosystem supporting AI innovation rests in the hands of a narrow group of investors whose values are not necessarily aligned with broader social needs.^[2] This private concentration of AI capability and wealth is structurally subsidized by the Department of Defense procurement of private-sector AI defense technologies as well as NSF-subsidized university training of future industry research scientists.^[3] In its most recent report, even the National Security Commission on Artificial Intelligence warns of the concentration of AI development in the hands of a few players and explicitly calls out the overwhelmingly commercial agenda dictating contemporary machine learning research.^[4] This current uneven state of AI access is not an accident, but the direct consequence of specific past policy choices determining what type of innovation research is taking place—and who benefits from it.^[5]

To extend the reach of AI innovation outside the small set of organizations already handsomely accruing benefits from the present system, a diversified approach to federal research and innovation funding that looks beyond the existing tactics of basic science funding, defense spending, and industrial commercialization programs is required. The federal government should play a larger role in funding AI research, design, and evaluation efforts aimed at addressing the needs of community-serving organizations such as hospitals, municipal governments, and health and human service providers. Translating advances in AI techniques into usable system designs for community-serving organizations requires its own form of research that community-serving organizations are not equipped to undertake on their own, and university researchers are not currently incentivized to pursue. Federal funding of community-centered AI research can therefore serve as a key pillar for AI innovation policy that helps seed a future in which the benefits of AI are more broadly distributed and in which the development of AI evolves in tandem with local needs and regional priorities.

AI Policy as an Arms Race

To date, American policy on AI has been primarily cast within the framework of national defense. The most important statutory provisions related to contemporary AI policy come from the National Defense Authorization Act of 2021 (NDAA). Among other things, the NDAA has established a new National Artificial Intelligence Initiative Office within the White House, has tasked the National Institute of Standards and Technology (NIST) to

develop collaborative frameworks, standards, and guidelines for AI, and has also ordered the Pentagon to ensure that the AI technologies it procures are developed in an ethical manner.^[6]

The military lens applied to contemporary AI policy discussions at the federal level is magnified by arguments that conjure the future of AI development in the language of an AI arms race.^[7] In this line of policy debate, of pressing concern is the risk that China—through its massive state-based tech-investment schemes—poses in demoting the United States from its dominant leading position in AI innovation on the global stage. AI as a matter primarily of national defense is echoed throughout the most recent report published by the National Security Commission on Artificial Intelligence (led by former Google CEO Eric Schmidt) which is organized around two themes: “Defending America in the AI Era” and “Winning the Technology Competition.”^[8]

Not only has the arms race framing served to justify increases in AI defense spending, it has also been leveraged successfully to lobby for increased funding of basic research in AI through the National Science Foundation (NSF).^[9] This is a familiar reflex, dating back at least to World War II, in which large-scale U.S. government support of theoretical research is triggered in response to external threats.^[10] While the research funded by NSF’s new National AI Research Institutes (NAIRI) arm holds some promise in expanding the impact of AI into selected civilian domains such as healthcare and agriculture, its overwhelming focus remains on basic theoretical research that is “use-inspired” but not actually application-oriented. As communicated to potential NAIRI applicants, “it is not the intent of the program that Institutes should focus mainly on the application of AI.”^[11] In aggregate, the current U.S. AI innovation research policy deliberately deprioritizes applied research, which NSF defines as “original investigation undertaken to acquire new knowledge; directed primarily, however, toward a specific, practical aim or objective.”^[12] As such, even research funding programs like NAIRI—ostensibly designed to better address the range of complex societal impacts engendered by AI technologies—doubles down on an existing bias favoring research that is disconnected from and unaccountable to attested societal needs or real-world domain practices. Currently, only 13% of NSF’s research funding goes to applied research.^[13]

Research on Integration Challenges is Key to AI Innovation

The development of AI systems on the ground is often stymied by a myriad of real-world adoption and integration problems that can only be understood and mitigated through rigorous applied research.^[14] AI systems are difficult to understand, rendering it difficult for

human operators to assess their safety or reliability.^[15] AI systems pose significant privacy concerns given the type and quantity of data required for effective model development.^[16] And AI models have been shown in many contexts to perform with less accuracy on minority groups, thus exacerbating existing injustices.^[17] This set of application-focused concerns has prompted the formation of successful civil society and grassroots campaigns leading to moratoria, and even outright bans, on some uses of AI technologies across major American cities.^[18]

Perhaps most importantly, AI systems are often implemented without regard for community needs or objectives. These disconnects may arise when AI systems designed for one context are ported over to a new context without attention to the fact that predictive models built for one environment may not perform as well in a new locality, and that organizations naturally differ in how they make use of AI systems.^[19] Designing and integrating AI systems for real-world use requires not only technical expertise but also deep domain knowledge and experience, as well as systems design and social scientific approaches to effectively bridge between the social and technical. Yet, while advances in the basic research underpinning AI continue to push forward the state-of-the-art, developing capabilities around the design and integration of AI systems has lagged behind.^[20]

AI innovation policy must be more inclusive in determining what types of innovation research receive funding to bridge this chasm. In particular, substantially more research funding should be dedicated to studying whether and how to translate and integrate the latest technical advances in AI into effective, beneficial applications that address concrete problems identified at the community level. Specifically, federal research agencies should develop AI programs that incentivize and forge cross-sector and cross-disciplinary research teams whose members are composed of domain experts, community stakeholders, and technical experts, and who together collaboratively undertake applied AI systems research, design, and evaluation. There are (at the time of publication) various congressional and executive proposals that seek to expand the scope and budget of the NSF substantially.^[21] While these proposals appropriately push the NSF to develop a new directorate focused specifically on technology and innovation concerns, these proposals all share a continued primary focus on fundamental research, with a secondary emphasis on commercialization research. An expanded, more robust NSF, however, must be reconfigured to foster translational research for technology in the public interest whose benefits are not constrained by commercial or defense goals.

Centering Community Needs within AI Innovation Policy

An AI innovation policy approach centered on community is not purely a theoretical counter to the current state of affairs. Across American society and economy, AI could play a central role in helping solve deep structural problems in the public interest. Promising results from pioneering applied research efforts conducted within local organizations demonstrate that AI can be harnessed in hospitals to help reduce clinician errors that lead to misdiagnosis,^[22] in municipal government to protect vulnerable residents from abuse by landlords,^[23] and in human services to enhance child welfare call screening decision-making.^[24] Additionally, researchers have begun investigating how to use AI within specific community contexts to help facilitate complex local coordination problems, such as ensuring the equitable and efficient distribution of food donations.^[25] Yet, for any area in which there are promising AI applications, there also exist a set of potential pitfalls and growing skepticism regarding these applications' economic value, trustworthiness, and compatibility with social and economic justice aims.^[26]

Addressing these integration challenges for AI merits its own dedicated research effort that goes beyond the received notion of applied research in traditional science and engineering policy.^[27] Applied research in an AI context requires cross-disciplinary examination of the design, fit, and maintenance of systems leveraging algorithms for particular problems, users, and communities.^[28] Law and medicine offer pioneering examples of this cross-disciplinary integrative approach to AI research that has led to the development of robust systems capable of serving mission-critical clinical decision-making functions.^[29] These efforts emphasized collaborative experimentation, learning, and deliberation across researchers, domain experts, and community stakeholders in order to develop systems that met the specific needs of the targeted organizations and the populations they serve. In achieving this goal, these efforts also generated baseline knowledge and lessons that equipped a new cohort of technology practitioners and scientific researchers to maintain and evolve the state of applied AI in the specific domain. These real-world precedents can be leveraged to devise a new type of AI innovation funding program that fosters the careful design and evaluation of AI within organizations and communities that are in need of novel solutions yet have limited AI access due to lack of resources and coordination.

A Proposal for Community-Centered AI Research Policy

Elsewhere, one of the authors of this article has [proposed](#) that the Biden-Harris administration adopt a program to foster community-centered AI research.^[30] This program consists of a direct investment into local research and development (R&D) efforts helping bring AI innovation capacity to American organizations that otherwise would not be able to shoulder its costs and complexity, while also incentivizing leading AI researchers to better

understand how to responsibly and effectively integrate AI into complex real-world situations. Through a comprehensive and multi-stakeholder research process, the projects funded by this program would serve as a catalyst to educate and empower practitioners and community stakeholders on the ground to take ownership of AI tools and processes, equipping them with the experience to carry forward the maintenance and evolution of applied AI practice in their domains. Additionally, these projects will serve to generate a valuable set of case studies of successes and failures, critical for developing a theory of AI integration theory and practice sorely missing in discussions of AI advances.

It is imperative to develop an AI innovation agenda that helps us better understand concretely how AI can be best applied in our current institutions to address contemporary problems.^[31] Received notions separating science and engineering research into discrete binary categories—fundamental vs. applied, theoretical vs. practical, experimental vs. implementational—all fail to address the needs of a contemporary moment in which technological advances do not (if they ever did) neatly stay within the confines of the lab before being set out for use in the larger population. Along these lines, the appointment of Eric S. Lander and Alondra Nelson to lead the White House Office on Science and Technology Policy (OSTP) is an encouraging development. In her remarks upon accepting the position, Dr. Nelson communicated the understanding that AI research and technology, like any science, is “at its core a social phenomenon. It is a reflection of people, of our relationships, and our institutions.”^[32] AI innovation policy needs to evolve to reflect this integrated view of what scientific knowledge is in order to engender the substantive and responsible societal impact to which it aspires.

The authors acknowledge support for their research from the Russell Sage Foundation, the John D. and Catherine T. MacArthur Foundation, and the Day One Project.

Fernando Delgado is a PhD student in [Information Science](#) at Cornell University whose research focuses on algorithmic system design and governance. Prior to commencing his doctoral studies, Fernando worked at [H5](#), a pioneering firm in the field of legal technology designing and deploying text classification systems for automating civil discovery review and fact-finding.

[Karen Levy](#) is an assistant professor in the Department of Information Science at Cornell University. Her research focuses on the social, ethical, and legal dimensions of data-intensive technologies.

1. Tom Simonite, “The Dark Side of Big Tech’s Funding for AI Research,” *Wired*, December 10, 2020; Daniel Zhang, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, Yoav Shoham, Jack Clark, and Raymond Perrault, *The AI Index 2021 Annual Report* (Palo Alto: Stanford Institute Human-Centered Artificial Intelligence, 2020). [↑](#)
2. Josh Lerner and Ramana Nando, “Venture Capital’s Role in Financing Innovation: What We Know and How Much We Still Need to Learn,” *Journal of Economic Perspectives*, Volume 34, Number 3, Summer 2020, pp. 237-261. [↑](#)
3. Justin Doubleday, “New analysis finds Pentagon annual spending on AI contracts has grown to \$1.4B,” *Inside Defense*, September 24, 2020; Dani Rodrik, “Democratizing Innovation,” *Project Syndicate*, August 11, 2020. [↑](#)
4. National Security Commission on Artificial Intelligence (NSCAI): *The Final Report*, 2021. <https://www.nscai.gov/2021-final-report/>. [↑](#)
5. Daron Acemoglu and Pascual Restrepo, “The wrong kind of AI? Artificial intelligence and the future of labour demand,” *Cambridge Journal of Regions, Economy and Society*, 2020, 13, 25-35. [↑](#)
6. William M. Thornberry *National Defense Authorization Act for Fiscal Year 2021. Conference Report to accompany H.R. 6395. 116th Congress (2020)*. [↑](#)
7. Graham Allison and Eric Schmidt, *Is China Beating the U.S. to AI Supremacy?* (Cambridge: Harvard Belfer Center for Science and International Affairs, 2020). [↑](#)
8. NSCAI (2021). [↑](#)
9. *National Artificial Intelligence (AI) Research Institutes: Accelerating Research, Transforming Society, and Growing the American Workforce*, Program Solicitation NSF 20-604, National Science Foundation, 2020. <https://www.nsf.gov/pubs/2020/nsf20604/nsf20604.pdf>; Tim Hwang, *Shaping the Terrain of AI Competition* (Washington, D.C.: Georgetown University: Center for Security and Emerging Technology, 2020); Remco Zwetsloot, Helen Toner, and Jeffrey Ding, “Beyond the AI Arms Race: America, China, and the Dangers of Zero-Sum Thinking,” *Foreign Affairs*, November 16, 2018. [↑](#)
10. Kei Koizumi, “The Evolution of Public Funding of Science in the United States From World War II to the Present,” *Oxford University Press and the American Institute of Physics*, March 31, 2020. [↑](#)
11. *Frequently Asked Questions (FAQs) about the National Artificial Intelligence (AI) Research Institutes Program (NSF 20-604)*, National Science Foundation, 2020, <https://www.nsf.gov/pubs/2020/nsf20123/nsf20123.jsp#q6>. [↑](#)
12. *Federal Research and Development(R&D) Funding: FY2020*, Congressional Research Service (R45715), March 2020, p. 2. <https://fas.org/sgp/crs/misc/R45715.pdf> [↑](#)

13. Federal Research and Development(R&D) Funding: FY2020 (March 2020, p. 37). [↑](#)
14. Sam Ransbotham, Shervin Khodabandeh, Kavid Kiron, François Candelon, Michael Chu, and Burt LaFountain, “Expanding AI’s Impact Organizational Learning,” MIT Sloan Management Review and Boston Consulting Group, October 2020. [↑](#)
15. Danielle C. Tarraf, William Shelton, Edward Parker, Brien Alkire, Diana Gehlhaus, Justin Grana, Alexis Levedahl, Jasmin Leveille, Jared Mondschein, James Ryseff, Ali Wyne, Dan Elinoff, Edward Geist, Benjamin N. Harris, Eric Hui, Cedric Kenney, Sydne Newberry, Chandler Sachs, Peter Schirmer, Danielle Schlang, Victoria M. Smith, Abbie Tingstad, Padmaja Vedula, and Kristin Warren, The Department of Defense Posture for Artificial Intelligence: Assessment and Recommendations (Santa Monica, CA: RAND Corporation, 2019). [↑](#)
16. Cameron F. Kerry, Protecting Privacy in an AI-driven World (Washington D.C: The Brookings Institution, Center for Technology Innovation, 2020). [↑](#)
17. Joy Buolamwini and Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” Proceedings of ACM Conference on Fairness, Accountability, and Transparency (2018): 77-91; Clare Garvie, Alvaro Bedoya, and Jonathon Frankle, The Perpetual Line-Up: Unregulated Police Face Recognition in America (Washington D.C: Center on Privacy & Technology, Georgetown Law, 2016). [↑](#)
18. Sidney Fussel, “The Next Target for a Facial Recognition Ban? New York,” Wired, January, 28, 2021. [↑](#)
19. Daniel E. Ho, Emily Black, Maneesha Agrawala, and Fei-Fei Li, Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains (Palo Alto: Stanford Institute Human-Centered Artificial Intelligence, 2020). [↑](#)
20. Tara Balakrishnan, Michael Chui, Bryce Hall, Nicolaus Henke, “Global Survey: The State of AI in 2020,” McKinsey and Company, November 2020; Kate Crawford, There is a Blind Spot in AI Research, Nature 538 (7625): 311-313, 2016; Kate Crawford, Roel Dobbe, Theodora Dryer, Genevieve Fried, Ben Green, Elizabeth Kaziunas, Amba Kak, Varoon Mathur, Erin McElroy, Andrea Nill Sánchez, Deborah Raji, Joy Lisi Rankin, Rashida Richardson, Jason Schultz, Sarah Myers West, and Meredith Whittaker, AI Now 2019 Report (New York: AI Now Institute, 2019). [↑](#)
21. U.S. Congress, Senate, Endless Frontier Act, S.3832, 116th Congress, <https://www.congress.gov/bill/116th-congress/senate-bill/3832/text>; U.S. Congress, House, National Science Foundation for the Future Act, H.R. 2225, 117th Congress, https://science.house.gov/imo/media/doc/NSF-FORTHEFUTURE_01_xml.pdf; Executive Office of the President, Office of Management and Budget, President’s Request for FY2022 Discretionary Funding, <https://www.whitehouse.gov/wp-content/uploads/2021/04/FY2022-Discretionary-Reque>

- st.pdf. [↑](#)
22. Nan Wu, Jason Phang, Jungkyu Park, Yiqiu Shen, Zhe Huang, Masha Zorin, Stanisław Jastrzębski et al. “Deep neural networks improve radiologists’ performance in breast cancer screening,” *IEEE Transactions on Medical Imaging* 39, no. 4 (2019): 1184-1194. [↑](#)
 23. Teng Ye, Rebecca Johnson, Samantha Fu, Jerica Copeny, Bridgit Donnelly, Alex Freeman, Mirian Lima, Joe Walsh, and Rayid Ghani, “Using Machine Learning to Help Vulnerable Tenants in New York City,” in *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies* (2019): 248-258. [↑](#)
 24. Alexandra Chouldechova, Diana Benavides Prado, Oleksandr Fialko, Emily Putnam-Hornstein, and Rhema Vaithianathan, “A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions,” *Proceedings of Machine Learning Research* 81:1-15, 2018. [↑](#)
 25. Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D. Procaccia, “WeBuildAI: Participatory Framework for Algorithmic Governance,” *Proceedings of the 2019 Conference on Human Computer Interaction, CSCW*, Article 181, 35 pages. [↑](#)
 26. Mariano Florentino-Cuéllar, David Freeman Engstrom, Daniel E. Ho and Catherine Sharkey, *AI’s Promise and Peril for the U.S. Government* (Palo Alto: Stanford Institute for Human-Centered Artificial Intelligence, 2020); Madeleine Elish and Elizabeth Watkins, *Repairing Innovation: A Study of Integrating AI in Clinical Care* (New York: Data & Society, 2020). [↑](#)
 27. Bammer, G., O’Rourke, M., O’Connell, D. et al, “Expertise in research integration and implementation for tackling complex problems: when is it needed, where can it be found and how can it be strengthened?,” *Palgrave Communications* 6, 5 (2020). [↑](#)
 28. Ben Green and Salomé Viljoen, “Algorithmic realism: expanding the boundaries of algorithmic thought.” *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAcT ’20)*. Association for Computing Machinery, New York, NY, USA, 19-31; Andrew Selbst et al. (2019). *Fairness and Abstraction in Sociotechnical Systems*. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAcT ’19)*. ACM, New York, NY, USA, 59-68. [↑](#)
 29. Fernando A. Delgado, “Sociotechnical Design in Legal Algorithmic Decision-Making,” *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing* (2020): 111-115; Mark Sendak, Madeleine Elish, Michael Gao, Joseph Futoma, William Ratliff, Marshall Nichols, Armando Bedoya, Suresh Balu, and Cara O’Brien, “The Human Body is a Black Box: Supporting Clinical Decision-Making with Deep Learning,” In *Proceedings of ACM Conference on Fairness,*

Accountability, and Transparency (2020): 27-30. [↑](#)

30. Fernando A. Delgado, A National Program for Building Artificial Intelligence within Communities, Day One Project, Federation of American Scientists, January 2021. [↑](#)
31. Diana Nucera, Berhan Taye, Sasha Costanza-Chock, Micah Sifry, and Matt Stempeck. Pathways Through the Portal: A Field Scan of Emerging Technologies in the Public Interest. NY, NY: Civic Hall, 2020. Available at <https://emtechpathways.org>. [↑](#)
32. ABC News. Biden picks Alondra Nelson as deputy science policy chief. <https://abcnews.go.com/US/video/biden-picks-alondra-nelson-deputy-science-policy-chief-75299191>. [↑](#)