

Partnership on Al response to Request for Information on the Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Background

Partnership on AI (PAI) is a non-profit partnership of academic, civil society, industry, and media organizations creating solutions to ensure that AI advances positive outcomes for people and society. PAI studies and formulates sociotechnical approaches aimed at achieving the responsible development and deployment of artificial intelligence (AI) technologies so that they are trustworthy, reliable, secure, and accountable. Today, we connect nearly 130 partner organizations in 15 countries to be a uniting force around this aim.

PAI develops tools, recommendations, and other resources by inviting multistakeholder voices from across the AI community and beyond to share insights that can be synthesized into actionable guidance. We then work to promote adoption in practice, inform public policy, and advance public understanding. We are not an industry or trade group nor an advocacy organization. We aim to change practice, inform policy, and advance understanding.

The information in this document is provided by PAI and is not intended to reflect the view of any particular Partner organization of PAI. The comments provided herein are intended to provide evidence-based information, based on PAI's research, in response to the Request for Information (RFI) on the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan.

Executive Summary

PAI recommends that the 2025 National AI R&D Strategic Plan be built on the 2023, 2019 and 2016 National AI R&D Strategic Plans, including by substantially retaining the 9 overarching Strategies listed in the 2023 Plan with discrete updates to account for new developments in the field of AI. There are 3 strategic reasons why building on these prior plans will position the 2025 National AI R&D Strategic Plan and the US most strongly:

- They reflect significant bipartisan agreement around strategic R&D priorities;
- 2. They strike a good balance between levering Al's tremendous opportunities and addressing its challenges;
- 3. They are grounded in multidisciplinary research and are backed multistakeholder support across industry, academia and civil society.

Building on our existing and ongoing research, we provide specific reasons for maintaining and expanding Strategies 3, 4, 5, 6, 7 and 9 in the sections below.

Summary of recommendations

Primary recommendation:

 The updated National AI R&D Strategic Plan should continue to be built on the foundation of the 2023, 2019 and 2016 Plans. The 9 Strategies in the 2023 National AI R&D Strategic Plan should be retained and updated consistently with this submission.

Overarching recommendations:

- 2. Both technical and sociotechnical research should be funded under the updated Plan.
- 3. The updated Plan should prioritize multistakeholder research.
- 4. The role of, and funding for, the NSF should be emphasized in the updated Plan.

Specific recommendations to strengthen the Strategic Plan (by expanding on or introducing new research areas):

- 5. Research to understand the ethical, legal, and societal implications of Al warrants continued Federal support (**Strategy 3**).
- 6. Research on the robustness and reliability of Agentic AI should be explicitly included as a priority (**Strategy 4**).
- 7. Research on the challenges and opportunities of open models should be included, alongside open datasets (**Strategy 5**).

- 8. Al assurance should feature more prominently in the 2025 National AI R&D Strategic Plan, including support for research on different assurance mechanisms, such as evaluations, benchmarks, testing, validation, verification, audits, and certification (**Strategy 6**).
- 9. Al's impact on the labor market and the broader economy should be included as a distinct research priority under Strategy 7. This should include research on job quality. The Plan should also be updated to address the linkages between the social and ethical impacts of AI on the labor force included under Strategy 3 with the AI workforce needs included under Strategy 7.
- **10. Strategy 9** should be expanded to include research on technical and policy interoperability across jurisdictions.

Overarching recommendations

PAI welcomes the opportunity to submit its recommendations to the Office of Science and Technology Policy (OSTP) and the National Science Foundation (NSF) on the 2025 National AI R&D Strategic Plan. We welcome OSTP's move to update the 2023 National AI R&D Strategic Plan given recent developments in the field of AI, including advanced reasoning capabilities, the emergence of and growing interest in AI agents and AI companions, efficiency gains, particularly by focusing resources on AI inference, and the growing integration of AI in daily life.

We believe that the 2023 National AI R&D Strategic Plan – which consistently builds on the 2019 and 2016 Plans – provides a good basis for the 2025 National AI R&D Strategic Plan, and its core elements should be retained. Together these plans have fostered US AI leadership across multiple administrations. Specifically, retaining the 9 overarching Strategies laid down in the 2023 Plan will ensure that Federal investments in AI R&D continue to leverage the opportunities of new and existing AI technologies while addressing their challenges. At the same time, we believe that specific updates should be made to some of those Strategies to take account of new developments in the field of AI, in line with the specific recommendations below.

As recognized consistently in the 2023, 2019 and 2016 National AI R&D Strategic Plans, designing AI and understanding its impact on individuals and society is a complex endeavor that cuts across different disciplines in the hard sciences (e.g., computer science, engineering, mathematics and biology), social sciences (e.g., law, economics, linguistics, psychology and anthropology) and humanities (e.g., ethics and philosophy). On this basis, **it is crucial that both technical and sociotechnical research across the 9 Strategies continue to be supported by the Federal government in the 2025 National AI R&D Strategic Plan.** In particular, the Federal government should continue to support programs and grants within the <u>NSF's Directorate for Social, Behavioral and Economic Sciences</u> (<u>SBE</u>), while broadening their remit to include AI-focused social and economic research.

In the same vein, to be holistic, AI R&D requires input from the various stakeholders that use and/or are impacted by AI technologies, including representatives of academia, industry, government and civil society. For example, as was highlighted in the <u>2019 National AI R&D Strategic Plan</u>, the NSF and PAI – a diverse, multistakeholder organization working to better understand AI's impacts – worked to jointly <u>support</u> high-risk, high-reward research at the intersection of the social and technical dimensions of AI.

Likewise, the engagement of the National Institute of Standards and Technology (NIST) with <u>multiple stakeholders</u> has been critical in the development of the AI Risk Management Framework (RMF). The RMF has been emulated by different jurisdictions around the world and has helped advance US leadership in the field. Therefore, like the 2023, 2019 and 2016 National AI R&D Strategic Plans, **the 2025 Plan should continue to recognize the need to support multistakeholder engagement in AI R&D.** Given the limited availability of Federal resources, the involvement of multiple stakeholders in a proposed project – particularly from academia, civil society and industry – should be a factor given weight in assessing applications for Federal research funding.

Crucially, PAI recommends continued support for the NSF, given its unique role funding scientific research that would not otherwise be carried out. As noted across the 2023, 2019 and 2016 National AI R&D Strategic Plans, rigorous academic research and inquiry are fundamental drivers of economic progress and human flourishing. **The 2025 National AI R&D Strategic Plan should explicitly recognize the importance of the NSF, including its work on open science**. In particular, the 2025 plan should support the NSF-led <u>National AI Research</u> <u>Institutes</u> and provide it with the necessary resources to implement the <u>National</u> <u>Artificial Intelligence Research Resource (NAIRR) Pilot</u> – both of which were launched under the National AI Initiative Act (2020). To make the most of limited Federal resources, the NSF should expand industry partnerships. This includes strengthening the role of the <u>Technology</u>. Innovation and Partnerships (TIP) in AI research, establishing AI-specific <u>Industry-University Cooperative Research</u> <u>Centers</u>, and creating industry-funded internships for graduate and undergraduate students.

Specific recommendations

Drawing on our existing and ongoing research, PAI makes the following specific recommendations on a number of the R&D Strategies included in the 2023 National AI R&D Strategic Plan, including reasons for their retention in the 2025 Plan and new proposals on how they should be updated. Recommendations for new or expanded research areas include: **agentic AI**, **open foundation models**, **AI Assurance, AI labor impacts**, and **technical and policy interoperability**.

Strategy 3: Understand and address the ethical, legal, and societal implications of AI

PAI has carried out extensive research and developed guidelines and recommendations on many of the impacts of AI in society, including, in particular, on <u>labor and the economy</u> and on the <u>information ecosystem</u>. Our work demonstrates that this line of research is crucial to ensure that AI has positive impacts on people and society: while AI can have incredibly beneficial applications – such as in healthcare, access to justice, and scientific discovery – it can also be exploited by malicious actors or inadvertently cause harm to different stakeholders. Addressing these risks is critical not just to prevent harm, but to build trust in AI systems. That is essential to drive adoption, which is key to US AI competitiveness. Numerous research gaps remain in this field that need to be addressed through sustained R&D investments. **As discussed further below, one key issue identified in our work is that research on labor market impacts should address job quality impacts in addition to job availability and wage impacts.**

We recommend that the Federal Government continue to invest in research to understand the ethical, legal, and societal implications of AI, as outlined in Strategy 3 of the 2023, 2019 and 2016 National AI R&D Strategic Plans. In particular, we recommend the creation of specific programs and grants on those topics, including within the NSF's SBE.

Strategy 4: Ensure the safety and security of AI systems

PAI has published <u>research</u> and recommended best practices to support key aspects of the trustworthy design and development of AI. These include <u>guidance</u> for foundation model deployment, best practices across the AI value chain, research on explainable AI, and <u>Incident Database</u> that PAI co-founded.

Our work shows that, to foster the design and deployment of trustworthy, reliable, secure and accountable AI, **research needs to continue on the development of**

tools and techniques to prevent and mitigate against both malfunction and misuse of AI technologies. This need has been consistently recognized in Strategy 4 of the 2023, 2019 and 2016 National AI R&D Strategic Plans and is especially important as AI capabilities advance, adoption becomes more widespread, and the stakes of using AI get higher. For those reasons, we support the retention of Strategy 4 in the 2025 Plan.

Notably, we recommend continued support for the <u>NSF's AI-Ready Test Beds</u>

Initiative, aimed at enhancing and advancing the essential infrastructure needed for researchers to develop, test and refine responsible AI systems in real-world settings. Testbeds can give developers the confidence they need to place their products on the market and further innovate. Given the sensitive nature of Federal data, tests on certain AI systems, especially those used by the Federal Government, can only be carried out in Government sandboxes such as those established by the NSF. In the same vein, certain types of tests like red teaming can expose vulnerabilities that might give rise to national security risks and therefore ought to be carried out in Federal testbeds.

What needs to be different from the prior approach?

PAI is currently carrying out research on AI agents, including failure monitoring, as well as devising a broader research agenda to help policymakers better understand the challenges and opportunities of AI agents as they gradually integrate our economy. While AI agents are still in their early days, it is clear from the research carried so far by PAI and <u>other stakeholders</u> that their increased autonomy and ability to take action in virtual or physical environments can raise new challenges and amplify existing ones. Research on AI agents and agentic systems is fairly new and requires sustained Federal R&D investments to continue to grow. **Therefore, we believe AI Agents and Agentic AI Systems should be included as research priorities in Strategy 4**.

Relatedly, while the RFI mentions 'advances in AI systems capable of reasoning, adaptability, and robustness in dynamic environments' as a possible research priority, we believe this should explicitly include, as part of Strategy 4, research into the robustness and reliability of AI agents and agentic AI systems. Notably, the field of benchmarks and evaluations for AI agents requires further research. This is because existing methods used for non-agentic LLM systems are <u>inadequate</u> to measure key attributes of AI agents, including their cost and reliability.

Strategy 5: Develop Shared Public Datasets and Environments for AI Training and Testing

Open frameworks, datasets and tools – such as the transformer architecture and PyTorch – have been crucial in the development of leading generative AI models within the US, and therefore been critical for US competitiveness. PAI has <u>supported</u> the development of open AI ecosystems to promote innovation and ensure that the benefits of AI are shared widely across America. At the same time, there <u>remains</u> a significant degree of uncertainty about what it means for AI to be 'open source' or 'open', and what the pros and cons of open versus closed models are. In particular, more data is needed on the risks of open models, including safety, security, and/or privacy vulnerabilities. Only a few private companies have invested in open AI models and there is generally little economic incentive for industry to build such models. This is because the whole point of open models is that they are freely available. Once released, open models can be widely adapted, are hard to monitor, and cannot be disabled. Government support is therefore needed for research both to promote the development of open models and to identify and address their risks.

For those reasons, we support the retention of Strategy 5 in the 2025 National AI R&D Strategic Plan. Specifically, alongside the development of open-source software libraries and toolkits, we support the inclusion of research on open models themselves, including their challenges and opportunities. We also recommend continued Federal support for the NSF-led Pathways to Enable Open-Source Ecosystems and the Safety, Security, and Privacy for Open-Source Ecosystems (Safe-OSE) program.

Strategy 6: Measure and evaluate AI technologies through standards and benchmarks

PAI has done substantial work on <u>post-deployment documentation standards</u> and is currently carrying out research on the various components of the AI assurance ecosystem, including how they relate to one another and what roles different stakeholders play in their implementation. **Standards and metrics are key components of a robust AI assurance ecosystem that seeks to build public trust and confidence in AI to unlock AI innovation, adoption and economic competitiveness**.¹ Without standards and metrics, it is difficult to make decisions about AI's levels of trustworthiness, reliability, security and accountability; and ultimately to drive progress in the field.

¹ See Executive Order 13859 of February 11, 2019, 'Maintaining American Leadership in Artificial Intelligence'.

It is critical that the Federal Government continues to support this work as there is insufficient economic incentive for industry to comprehensively assure AI technologies. While many private companies deem it important to assess models and systems before, during and after release, they are often under immense pressure from their shareholders to launch their products as fast as possible. In addition, start-ups and smaller companies developing AI do not always have the resourcing to develop and conduct robust assessments in-house.

The need for research on standards and benchmarks was noted in Strategy 6 of the 2023, 2019 and 2016 National AI R&D Strategic Plans.² **The Federal Government should continue to fund research that supports the development of AI standards as well as different methods to measure AI technologies against those standards**. Research in this area should continue to cover different attributes of AI technologies, including their capabilities, reliability, explainability, interpretability, transparency, privacy and compliance with civil rights and liberties, such as freedom of speech and non-discrimination. **In particular, the Federal government should continue to support the work carried out by the NIST, including through the AI Safety Institute (AISI), the AISI Consortium and the Network of AI Safety Institutes.**

What needs to be different from the prior approach?

However, we believe that AI assurance should feature more prominently in the 2025 National AI R&D Strategic Plan, including support for research on different assurance mechanisms, such as evaluations, benchmarks, testing, validation, verification, audits, and certification. Open research questions in this area include how to develop effective and reproducible socio-technical evaluations for generative AI's wide range of use cases and avoiding model overfitting. AI assurance is a booming field of research and a growing global market. Federal support for this work can secure American leadership in this important area. It is particularly important to continue to support AISI's work, including through its Consortium, as well as to establish NSF-led programs and partnerships that support academic and multistakeholder research efforts on AI assurance.

Strategy 7: Better understand the national AI R&D workforce needs

The 2023, 2019 and 2016 National AI R&D Strategic Plans reflect the need to expand R&D to better understand the needs of American workers as they prepare for an AI-driven economy. PAI submits this focus should be retained

and expanded on in the updated Plan. Our <u>research</u> has shown that while AI has tremendous potential to solve major societal problems and make peoples' lives better, it will drive new forms of automation, wealth distribution, and economic decision-making that we as a society must grapple with. Key AI impacts will include both workforce disruptions and the creation of new jobs. We therefore support the retention of all research areas listed in Strategy 7 of the 2023, 2019 and 2016 Plans, with particular emphasis on the need for more data on the current AI workforce and building AI literacy skills across different demographics.

PAI has produced <u>guidelines</u> to help policymakers, industry and other stakeholders prepare for the impacts of AI on the labor market and the economy to ensure that AI promotes shared prosperity. But the extent of those impacts, particularly long-term impacts on employment and wages across different sectors of the economy, remains under-explored and requires comprehensive investigation — including more data on AI adoption. It is especially important to assess the implications of increasingly autonomous AI agents on the labor market and to design appropriate policy interventions. We therefore support the inclusion of AI's impact on labor demands and the economy more generally as a distinct research priority under Strategy 7.

What needs to be different from the prior approach?

PAI has carried out research on <u>worker well-being</u> and <u>data enrichment workers</u> specifically. Further research is needed in both areas. This includes the need to develop critical-thinking skills, safeguard other worker skills and autonomy, and foster more transparency on data supply chains. Poor working conditions may result in low quality data, and consequently unreliable AI or harmful models. **We believe these issues should also feature in a revised version of Strategy 7.**

The Plan should also be updated to address the linkages between the social and ethical impacts of AI on the labor force included under Strategy 3 with the AI workforce needs included under Strategy 7.

Strategy 9: Establish a Principled and Coordinated Approach to International Collaboration in Al Research

PAI submits that Strategy 9 of the 2023 National AI R&D Strategic Plan, addressing international collaboration in AI research, should be retained. The need to cooperate with like-minded international partners on AI R&D was previously recognized in the 2019 Plan under Strategy 8. PAI's <u>research</u> has identified several benefits of international cooperation in the field of AI. In particular, it can help address AI risks that transcend national borders, embed US values in the development of AI technologies, and avoid a leadership vacuum that could be filled by authoritarian States. For example, the US AI Safety Institute has successfully <u>collaborated</u> with the UK (now) Security Institute on the evaluation of Open AI's o1 model, pursuant to a <u>memorandum of understanding</u> between the two countries. As AI becomes more widespread around the world, cooperation with allies in AI research – especially in the areas of testing, evaluation, validation and verification (TEVV) – is key to protecting US national security from foreign threats.

For those reasons, we recommend continued Federal support for the <u>NSF's</u> <u>Office of International Science and Engineering (OISE)</u>. This includes expanding or establishing AI-specific programs or opportunities that seek to, among other things, facilitate international exchanges among researchers, establish bilateral or multilateral research centers with allies, and co-fund international research projects with international partner institutions, such as UK Research and Innovation (UKRI) and the Natural Sciences and Engineering Research Council of Canada (NSERC).

What needs to be different from the prior approach?

The 2025 National AI R&D Strategic Plan should drive forward research on interoperability across AI policy frameworks. Our work has noted how interoperability is key to avoid inconsistent requirements and insurmountable compliance burdens for AI developers and deployers in the US. But more work is needed on what requirements or best practices should be interoperable, such as <u>documentation</u>, testing and evaluations, and how to implement them among developers and deployers. We therefore support the expansion of Strategy 9 to include not only cooperation in AI research but also research on technical and policy interoperability across jurisdictions.

Conclusion

For further information or any questions related to this submission, please contact john@partnershiponai.org and talita@partnershiponai.org.