PUBLIC SUBMISSION

Received: May 09, 2025
Tracking No. mag-w0g3-2ffw
Comments Due: May 28, 2025 **Submission Type:** API

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0069
Comment on FR Doc # 2025-07332

Submitter Information

Name: sebastien laye

General Comment

Docket ID No. NSF-2025-OGC-0001
Filed individually but also on behalf of 2 institutions: The Economic Singularity institute (think tank) and Republicans Overseas.

The document lays out pure research themes (the first section), but thereafter we also recapped industrial/ infrastructure AI public policy proposals.

Attachments

National Artificial Intelligence Strategic Plan Sebastien Laye

National Artificial Intelligence R&D Strategic Plan (2025) – Policy Proposal

Author: Sebastien Laye, Economist & Entrepreneur specializing in the Economics of AI. On behalf of The Economic Singularity Institute and Republicans Overseas.

"This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution."

Introduction:
The United States must adopt an ambitious yet balanced roadmap to advance AI research and development, maintaining leadership in the face of fierce global competition. This policy proposal outlines a comprehensive strategy for the National Science Foundation's 2025 National AI R&D Strategic Plan. It emphasizes diversified technical pathways toward Artificial General Intelligence (AGI) and even future **Artificial Superintelligence (ASI)**, robust measures to **outpace strategic competitors like China**, and structural innovations in our national R&D ecosystem. Drawing on expert insights, the plan details actionable recommendations (grouped thematically) to secure U.S. leadership in safe and advanced AI development. These recommendations span fundamental research priorities, geopolitical and economic strategies, institutional reforms, and bold "moonshot" initiatives, culminating in the ten most critical actions for policymakers.

1. **Powerful AI or AGI via Diversified R&D Pathways**

To achieve transformative AI capabilities, the U.S. must pursue **multiple technical routes to AGI** rather than betting on a single paradigm. This hedges against uncertainties and maximizes our odds of success in developing human-level AI. Key research avenues include:

- **Sustain LLM Frontier Scaling (with Caution):** Continue advancing large-scale **large language models (LLMs)** and other deep networks as one path toward AGI, while recognizing current limitations. Even as top models demonstrate impressive feats, there are **empirical and theoretical impediments** to unlimited scaling. Federal support should drive **algorithmic innovations (e.g. retrieval, fine-tuning)** to improve reasoning and efficiency in next-generation models, not just brute-force scaling. Policymakers should plan for the possibility that LLM-focused approaches could yield AGI in the near term, **without over-optimizing on this outcome alone**.

- **Invest in Neuro-Symbolic and Modular AI:** Prioritize research combining neural networks with symbolic reasoning or algorithmic modules. **Neurosymbolic AI** merges data-driven learning with explicit knowledge representations, enabling logical reasoning and explainability beyond what pure LLMs can do. By integrating **symbolic logic and deep learning**, these hybrid systems aim to achieve more **human-like reasoning, common-sense understanding, and problem-solving** capacity. Leading experts estimate that achieving human-level reasoning will likely require a substantial **symbolic AI component (≥50%)** alongside learning algorithms. Federal programs should fund projects on neurosymbolic architectures, **reasoning-centric AI** (e.g. theorem-proving AI, program synthesis), and other **modular designs** that address LLMs' gaps in logic and transparency. This will expand the nation's options for reaching true general intelligence.

- **Develop World Models and Cognitive Architectures:** Support research into AI architectures that possess explicit **"world models"** – internal simulations of reality that allow for reasoning, prediction, and planning. Current transformer-based networks lack persistent memory and real understanding of the world. By contrast, approaches like **Yann LeCun's Joint Embedding Predictive Architecture (JEPA)** propose building AI with predictive world modeling and embedded knowledge, enabling an AI to learn how the world works rather than just statistical patterns. Advancing such architectures (e.g. systems with long-term memory stores, modular perception-reasoning-planning components, or hierarchical cognitive loops) can close critical gaps in **reasoning, planning, and context retention** that LLMs currently exhibit. This research area will facilitate AI that **learns continuously**, adapts to real-time feedback, and reasons through complex, multi-step tasks – all essential for human-level intelligence. The NSF should issue targeted grants for projects on memory-augmented neural nets, neuromorphic cognitive architectures, and **hybrid AI agents with "System 2" deliberation** (e.g. combining neural learning with explicit planning algorithms).

- **Advance Embodied and Situated AI:** Make **embodied cognition** a core focus of AGI research. Intelligence arises not only from data but from interacting with the physical world – a point emphasized by both cognitive science and the recent AAAI 2025 Presidential Panel. An agent confined to text or images "cannot learn how to make decisions and act for itself in the world". Thus, the Plan should fund R&D in robotics and **embedded AI systems** that learn through physical or simulated interaction. For example, an embodied agent (a robot) with rich sensors can explore its environment to gain the kind of **common-sense knowledge** and causal understanding that static LLM training cannot provide. Projects might include developmental robotics (AI that learns like a child through curiosity-driven exploration), **simulation environments for AI** (virtual worlds to train agents safely at scale), and integration of language models with robotics (to give robots high-level reasoning abilities). By grounding AI in the real world, we enable it to develop **spatial, physical, and causal cognition** that is critical for general intelligence. Notably, U.S. efforts here will also counter rival nations (like China) that are prioritizing "embodied AI" for strategic reasons.

- **Explore Multi-Agent AI Systems:** Encourage research into architectures where **multiple AI agents collaborate or compete**, as a path to more robust intelligence. Complex cognition may emerge from systems of interacting agents (each with specialized roles or knowledge). Importantly, **teams of AI agents can be used to enhance reliability**, with agents checking and correcting each other's outputs. This concept, suggested by experts to improve factual accuracy, can be expanded into a paradigm where diverse agents (planners, critics, solvers) collectively solve problems that a single monolithic AI could not. Federal programs should fund experiments in **agent-based systems**, including **autonomous AI swarms** for problem-solving and **cooperative AI frameworks** for improved decision-making. Such research aligns with recent findings that next-generation AI might involve "**cooperating teams of agents** that continually fact-check each other" to reach higher levels of trustworthiness and performance.

- **Pursue Meta-Learning and Self-Improving AI:** Initiate research on AI systems that can **learn how to learn** and even improve themselves over time. Rather than training a model once on static data, meta-learning approaches allow AI to accumulate knowledge and adapt to new tasks continuously. In the long run, a transformative goal is **self-improving AI** – AI that can autonomously refine its own algorithms or architecture. Early steps include **AutoML and neural architecture search** (using AI to design better AI models) and **reinforcement learning self-play** techniques (e.g. AlphaGo's strategy of iteratively training against itself). The U.S. should treat

this as a moonshot research area: for example, sponsoring challenges for AI that can modify its own code to solve novel problems or **optimize its performance without human intervention**. Such work will lay the groundwork for future ASI, wherein machines could iterate on improvements much faster than humans. Ensuring this is approached safely (with proper alignment, see Section 4) will be paramount.

- **Support Next-Generation Hardware (Neuromorphic & Beyond):** Expand research into **novel computing substrates for AI** that could leapfrog current limitations. Today's advanced AI runs on power-hungry GPU clusters, and scaling further is constrained by energy and supply-chain bottlenecks. **Neuromorphic computing** offers an alternative: chips inspired by the brain that use spiking neural networks to achieve vastly greater energy efficiency. U.S. agencies (NSF, DARPA) should invest in neuromorphic prototypes and pilot programs to apply them in AI research. In parallel, exploratory funding can be provided for **biological computing** (e.g. networks of cultured neurons that compute with biochemical energy) and quantum-assisted AI – any technology that might enable orders-of-magnitude more efficient or capable learning. Embracing these cutting-edge hardware paths will ensure America leads not just in AI algorithms, but in the physical platforms that power future AI. As RAND experts note, pushing such **alternative substrates** to industrial scale could **upend AI supply chains and energy usage** in our favor.

By diversifying R&D across these pathways, the U.S. would hedge against uncertainty in how and when AGI will emerge. This strategy recognizes that while **LLMs have shown great promise**, it is *far from guaranteed* that simply making them bigger will reach human-level intelligence. To remain at the forefront, U.S. policy must "**plan for uncertainty and accommodate multiple pathways to AGI**" rather than guessing the winning approach in advance. In practice, this means a broad portfolio of research investments – from algorithms to hardware – all oriented toward the long-term goal of safe, autonomous intelligence. Such a portfolio will strengthen our national AI ecosystem, spur cross-pollination of ideas (e.g. insights from robotics informing language AI and vice versa), and reduce the risk of strategic surprise. The recommendations above ensure that if AGI is achievable, the United States will lead in its realization, whichever path it takes.

2. Competing and Leading in the Global AI Race

The geopolitical stakes of AI leadership are enormous. The U.S. must treat AI advancement as a **strategic priority on par with other national security imperatives**, given AI's potential to reshape economic and military power. China, in particular, has declared its ambition to lead the world in AI by 2030 and is investing heavily to that end. This Strategic Plan calls for proactive measures to **out-compete adversaries** and secure U.S. dominance in key areas of AI:

- **Scale Up National AI Funding and Ambition:** Substantially increase federal funding for AI R&D to **at least rival Chinese levels** (which are backed by massive state investments and industrial policy). The U.S. should build on the CHIPS and Science Act by appropriating new funds specifically for AI research centers, fellowship programs, and moonshot projects. This includes fully resourcing the NSF's AI institutes and establishing new **AI research hubs** at national labs and universities. Clear **national goals** should be set – for example, achieving certain breakthroughs (in AGI, in trusted AI, in robotics) by specific target dates – to galvanize efforts akin to a "space race" for AI. America's innovation engine, powered by free enterprise and academia, can outperform centralized authoritarian models, but only if we mobilize it with adequate support and urgency.

- **Establish AI Compute Infrastructure as Critical National Capital:** Treat access to computing power as a strategic resource. Just as past eras raced to build railroads or electrical grids, the AI era demands **vast compute infrastructure** for training advanced models. We recommend creating **National AI Computing Centers** – geographically distributed facilities offering petaflop-to-exaflop scale computing accessible to researchers nationwide. In line with the vision of **Special Compute Zones**, the U.S. should leverage federal land and retiring industrial sites to host these data centers. This includes repurposing decommissioned coal and nuclear plant sites (which already have robust power and grid connections) as AI supercomputing hubs. To coordinate this build-out, the President should appoint an **"AI Infrastructure Czar" or give this power to the US Investment Accelerator,** empowered to cut through red tape and fast-track high-performance computing projects. These measures will ensure American researchers always have world-class compute resources at their disposal – a decisive advantage as model sizes and data needs grow exponentially.

- **Secure the AI Supply Chain and Talent Pipeline:** The U.S. must guarantee supply of the critical inputs to AI innovation – namely advanced semiconductors and human talent. On semiconductors, we should **reinforce export controls** to deny adversaries access to top-end AI chips, while bolstering domestic chip manufacturing (through CHIPS Act implementation and public-private partnerships) so that our AI efforts are not bottlenecked by foreign supply. This includes investing in next-gen chip R&D (e.g. AI accelerators, neuromorphic chips) to maintain a **technology edge**. Equally important is expanding programs to **attract and retain global AI talent**. The U.S. should streamline visas for AI researchers, expand scholarships in AI-related fields, and incentivize our brightest STEM students to enter AI and robotics. A national workforce initiative can train tens of thousands of AI engineers and technicians, ensuring we have the human capital to implement our R&D goals. If China is producing more PhDs in AI, we will counter by making America the most attractive destination for top researchers worldwide. Human talent is one arena where America's open society has a natural advantage – one we must continue to leverage.

- **Leverage Allies and Build AI Coalitions:** Competing with China doesn't mean going it alone. The U.S. should deepen collaboration with allies (EU, G7, India, Israel, etc.) on AI research and standard-setting. Joint research programs, shared computing facilities, and talent exchange can amplify our collective capabilities while depriving adversaries of easy gains. We should also lead in establishing **international norms and alliances** for AI (for example, expanding the Global Partnership on AI) to ensure democratic nations set the rules for AI use and governance. By aligning with like-minded partners, we create a unified front that magnifies the West's edge in innovation and limits the influence of autocratic AI paradigms (such as China's export of AI surveillance tech).

- **Monitor and Respond to Adversary Advances:** Finally, the U.S. must **closely track China's progress** and be prepared with responses. If China pours billions into a particular AI domain, the U.S. should evaluate a proportional or asymmetric investment. For instance, China recently designated "**embodied AI**" and humanoid robotics as priorities, backing them with a $138 billion state fund. The U.S. should not allow any singular technology domain (be it intelligent robots, AI-driven biotech, etc.) to become a one-sided Chinese advantage. This Plan recommends a standing interagency "AI Competitive Intelligence" unit to continuously assess global developments and inform leadership of needed strategic responses. We will win the AI race not by panic or imitation, but by confidently out-innovating our competitors at every turn.

3. Techno-Industrial Initiatives for AI Leadership

Maintaining U.S. supremacy in AI will require not just *what* we research, but *how* we organize and fund that research. We must update our institutions and incentives for the AI era, drawing on recommendations from the recent **Techno-Industrial Policy Playbook (2025)** and other forward-looking policy research. The following structural initiatives will empower breakthrough AI R&D and accelerate the transition of ideas from lab to real-world impact:

- **Launch the "X-Labs" Program for Transformative Research:** Create a new class of federally-funded research organizations – **X-Labs** – dedicated to high-risk, high-reward science in AI and related fields. Traditional grant funding (NSF, NIH) tends to be cautious and incremental, leaving a gap in support for ambitious, interdisciplinary projects. The X-Labs initiative (as proposed by policy experts) would fill this gap by providing **long-term, flexible funding to entire research teams/institutions** rather than individual short-term projects. We recommend establishing an X-Labs competitive program with several tiers of awards, for example: **X01 "Excellence" centers** funded ~$10–50M/year for up to 7 years to pursue fundamental breakthroughs (modeled after institutes like Janelia Research Campus or Allen Institute); **X02 "Execution" labs** focused on mission-oriented engineering challenges (e.g. developing an open-source GPT-level model or new AI training platform); and smaller X04 exploratory grants for radical ideas in nascent stages. Crucially, X-Labs would introduce *dynamism* into federal R&D – labs are funded for a term and must prove their impact to be renewed, with a hard cap to ensure rotation. By funding **20+ such labs**, the U.S. can create hotbeds of creative AI research insulated from short-term pressures, driving the kind of disruptive innovation needed for AGI and beyond.

- **Establish Special Compute Zones (SCZs):** In tandem with national computing centers (Section 2), the U.S. should designate **Special Compute Zones** that incentivize the clustering of AI infrastructure and talent in strategic locations. An SCZ would offer expedited permitting, tax advantages, and readily available power for companies and institutions building large data centers or AI testbeds. The Techno-Industrial Playbook outlines this concept in detail: such zones could be planned on federal lands or repurposed industrial sites where gigawatt-level power is accessible. For example, decommissioned coal plant sites (with existing transmission lines) or regions with abundant renewable energy could become SCZs for AI. Within these zones, an **AI infrastructure czar** (or dedicated agency office such as the recently announced US Investment Accelerator) should coordinate swift development – e.g. using Defense Production Act authorities to prioritize critical equipment and pre-clearing environmental reviews on "previously disturbed" lands to speed up construction. The goal is to **dramatically shorten timelines** for standing up the mega-clusters of compute that frontier AI research demands. By lowering infrastructure barriers, SCZs will attract both private and public AI labs, creating regional innovation hubs that keep the U.S. at the cutting edge.

- **Revitalize Funding Mechanisms for Innovation:** Beyond X-Labs, other funding reforms are needed to unleash American innovation. The NSF and other agencies should expand support for **Focused Research Organizations (FROs)** – time-limited, goal-driven teams addressing specific infrastructure or data challenges in AI (akin to the X02 concept). Likewise, the **Small Business Innovation Research (SBIR)** program should be retooled to better support AI startups, with faster turnaround and greater Phase II/III funding for scaling successful prototypes. We also propose experimenting with **novel grant models** in civilian R&D: for instance, ARPA-style milestone payments in academic grants, or "lottery" funding pools for high-risk ideas (to reduce proposal overhead). Another idea is to devote a portion of NSF/DOE funds to *long-shot* proposals each year – acknowledging that even failures provide learning. The common aim is to **increase funding flexibility and risk-tolerance** in the federal portfolio, so that transformative

ideas in AI (which often cut across disciplines and may seem unconventional) have a better chance of obtaining support. Our competitors are pursuing breakthrough technologies with aggressive state-led programs; the U.S. must respond by empowering its innovators through smarter funding approaches.

- **Strengthen Public-Private Partnerships and Tech Translation:** The federal government should actively partner with industry leaders and entrepreneurial entities to accelerate AI R&D. This can take many forms: **co-funding arrangements** where government provides base research funds and industry contributes compute resources or data; joint AI testbeds (e.g. a government cloud platform where academia, startups, and companies collaborate on large experiments); and commercialization programs that help transition research out of labs. For example, NSF could expand its **Industry-University Cooperative Research Centers** focusing on AI, and DARPA could launch challenges in areas like reliable AI or edge AI with cost-sharing from tech companies. The Plan also supports creating **regional innovation ecosystems** (as in the CHIPS Act) for AI – linking universities, local industry, venture capital, and government support in key hubs (e.g. Silicon Valley, Boston, Austin, Pittsburgh). A vibrant public-private approach leverages the strengths of both sectors: academia's long-term vision and industry's ability to scale. It will speed up the **translation of breakthroughs into real products and services**, reinforcing U.S. economic leadership in AI. Notably, such collaboration is already yielding results – OpenAI's recent "AI jam sessions" with national labs to apply new reasoning models is a prime example of synergy between government research and cutting-edge AI firms. We must encourage more of this cross-pollination.

- **Promote Data and Resource Sharing Initiatives:** The quality and scope of AI research is often limited by access to large datasets and computing resources, especially for academic and non-profit researchers. A national strategy for **open datasets and compute access** can multiply our R&D output. This includes funding the creation of high-value **open datasets** in areas like science, healthcare, and defense (while respecting privacy and security), and incentivizing companies to contribute anonymized data for public research use. Additionally, expanding programs like NSF's CloudBank which provide cloud credits to researchers, or establishing a **national research cloud**, will level the playing field so that university teams can experiment at scales comparable to industry. Timely access to data is also a strategic issue – as highlighted in the Techno-Industrial Playbook's "securing foreign data flows" chapter, U.S. policymakers should ensure that American AI developers can obtain diverse global datasets (through data alliances or negotiated agreements) to train the best models. In summary, we must treat data as the new oil and computing power as the new steel of the modern economy – and adopt policies that ensure American researchers have abundant access to both.

These institutional and policy innovations create a **fertile environment for AI breakthroughs**. The X-Labs model directly addresses structural barriers in our R&D system – namely, the difficulty of sustaining **large-scale, interdisciplinary, and long-term research** within current funding mechanisms. By providing generous, stable support to top talent pursuing bold ideas, we increase the odds of game-changing discoveries (much as past government initiatives did in the Cold War era). Special Compute Zones and improved infrastructure planning recognize that **AI advancement is as much an engineering challenge as a scientific one** – we need the physical capability to execute our ambitions, not just blueprints. Overall, embracing a techno-industrial mindset means using all the tools of governance (funding, regulation, convening power) to accelerate innovation. In an era when technological leadership determines economic prosperity and security, updating our national innovation playbook is

not optional – it's imperative. The above measures will help ensure the U.S. research community stays **agile, well-resourced, and mission-focused**, delivering the AI advances upon which our future hinges.

5. AI for Scientific Discovery and Innovation

One of the most exciting promises of advanced AI is its potential to **revolutionize scientific research itself**. By serving as intelligent assistants – or even autonomous scientists – AI systems can dramatically accelerate the pace of discovery across fields. The U.S. should harness this potential through initiatives that integrate AI into the scientific process, essentially amplifying our national innovation output. Key proposals could include:

- **Initiate an "AI Scientist" Program:** We propose a flagship NSF program to develop AI systems capable of **autonomous scientific research** – from generating hypotheses to designing and running experiments, and even making theoretical insights. Recent progress foreshadows this possibility: e.g. DeepMind's AlphaFold2 cracked the decades-old protein folding problem, contributing to a scientific breakthrough worthy of accolades. Experimental platforms like the robot scientists "Adam" and "Eve" have already autonomously identified new drug candidates in lab settings. Building on these successes, the AI Scientist program would fund teams to create AI-driven automated labs in areas like chemistry (for material or drug discovery), physics, and biology. The ultimate moonshot here is the **"Nobel Turing Challenge"** – a grand challenge proposed in the community that calls for an AI system to produce Nobel Prize-caliber scientific discoveries. By 2030, the U.S. should aim to demonstrate highly autonomous AI systems that can meaningfully contribute to research at a human-expert level. This will not only yield specific breakthroughs (in medicine, energy, etc.) but also cement U.S. leadership in the meta-field of AI-driven science.

- **Deploy AI Tools to Supercharge Human Researchers:** In the nearer term, we must deploy generative AI and machine learning tools widely to assist human scientists and engineers. This includes **AI-driven hypothesis generation** (finding patterns in literature and data to suggest new hypotheses), **automated data analysis** (using ML to find signals in experimental data), and **simulation and modeling** (AI to create more accurate models of complex systems, from climate to cosmology). OpenAI's "**O3**" and "**O4-mini**" reasoning models are examples aimed at aiding scientific idea generation by drawing on cross-disciplinary knowledge. We should partner with organizations like OpenAI, Google DeepMind, and top research universities to integrate such models into research workflows. For instance, national laboratories could host **AI assistance workshops** ("AI jam sessions") where domain experts team up with AI experts to apply latest models on tough research problems. We also endorse making state-of-the-art scientific AI tools (like AlphaFold or advanced weather prediction models) available through cloud platforms to any U.S. researcher. The payoff will be seen in faster discoveries – what used to take months of trial and error could be achieved in days with AI-augmented insight.

- **Create AI-Driven Innovation Labs in Key Sectors:** Establish specialized innovation labs that use AI to push forward strategic national priorities. For example, an **AI for Drug Discovery Lab** could integrate robotics for high-throughput screening with AI models that predict molecular behavior, aiming to slash the time and cost of developing new medications. An **AI for Climate and Energy Lab** might use AI to optimize fusion reactor designs or discover new battery materials. Funding for these labs can be joint between science agencies (NSF, NIH, DOE) and industry partners, ensuring practical focus. The startup ecosystem is already moving in this direction (e.g. Lila Sciences' plan for a "scientific superintelligence platform" with fully

autonomous labs shows private capital interest). By launching public-private institutes around such themes, we harness AI as a force multiplier on our most pressing scientific and industrial challenges. Each lab would serve as a proving ground for AI methodologies that could be transferred to other domains as well.

- **Leverage AI for Engineering and Design ("AI-Driven Engineering"):** Beyond pure science, AI can dramatically improve engineering R&D via generative design, automated coding, and optimization. We recommend initiatives where AI helps design complex systems – for instance, using AI to co-design new semiconductor architectures (important for maintaining Moore's Law), aerospace vehicles, or resilient infrastructure. **Generative AI models can propose thousands of design variations** under given constraints, from which engineers can select promising candidates. In software, AI coding assistants (like Codex) already boost productivity; future **self-debugging, self-optimizing code systems** could take this further. An NSF program on **AI-augmented engineering** would spread these practices, funding collaborations between AI scientists and traditional engineering disciplines. The economic impact of widespread AI-assisted design and innovation could be enormous, increasing the competitiveness of U.S. manufacturing and technology firms. In essence, we aim to couple America's unrivaled engineering know-how with AI's speed and breadth, to out-innovate any competitor.

- **Empower Citizen Science and Education with AI:** A complementary aspect is making advanced AI tools available to students, educators, and citizen scientists. By integrating AI into STEM education (for example, AI tutors that can answer scientific questions or help with experiments), we train the next generation to effectively use these tools. Citizen science projects (such as crowdsourced data collection) could benefit from AI analysis or coordination. The NSF could sponsor challenges for high school or undergraduate teams to use AI in research projects, fostering early familiarity. Democratizing AI for discovery ensures the benefits extend beyond elite labs, and creates a culture where human creativity works hand-in-hand with AI. This will cultivate public support and understanding of AI as a positive force in advancing knowledge.

Using AI to accelerate science creates a virtuous cycle: AI helps make discoveries, which in turn can improve AI (for instance, discovering new algorithms or materials that improve computing hardware). The U.S. has always led through scientific innovation – AI offers a chance to supercharge this engine. Moreover, if we succeed in AI-enabled discovery, we secure leadership in both AI *and* the multitude of fields AI touches (from pharma to materials to clean energy). Other nations are beginning to explore this nexus – we must move faster and more systematically. By embedding AI deeply into our scientific enterprise, we ensure that **American researchers will solve problems faster and stay on the cutting edge**. A recent survey noted that a wide spectrum of approaches – from AI as a collaborative tool to fully autonomous AI scientists – will fundamentally transform how we do science. We embrace that transformation and intend to lead it. Ultimately, the breakthroughs enabled (a cure for a disease, a new energy source, etc.) will themselves justify the investment many times over, securing national prosperity and well-being.

Conclusion: Top 10 Strategic Recommendations

In summary, the following ten recommendations are the most critical, cross-cutting actions for the NSF's 2025 National AI R&D Strategic Plan. These represent our **highest priorities** to advance U.S. leadership in AI research and development:

1. Embrace a Multi-Pathway Approach to AGI: Commit to a diversified R&D portfolio exploring all plausible routes to advanced AI – from scaling deep learning to neurosymbolic hybrids, embodied agents, and beyond – rather than relying on one paradigm. This hedges our bets and maximizes the chance of breakthroughs. Policymakers should ensure funding and attention are distributed across these approaches, with mechanisms to rapidly support new promising avenues.

2. Establish National AI Computing Infrastructure & Special Zones: Treat AI compute as strategic infrastructure. Launch **Special Compute Zones** and national AI supercomputing centers to provide researchers with unparalleled computing power. By repurposing industrial sites and leveraging federal lands, rapidly build out the data centers needed for training tomorrow's frontier models, keeping America ahead in the compute race.

3. Launch X-Labs for Transformative AI Research: Create the X-Labs program to fund elite research institutes focused on long-term, high-risk AI science. By granting 5–7 year funding to top-tier teams (with flexible, large budgets), X-Labs will enable the kind of ambitious experimentation – in AGI algorithms, novel chips, etc. – that traditional grants cannot. This structural innovation will fill the gap in our R&D ecosystem and drive transformative progress.

4. Bolster AI Talent and Knowledge Networks: Expand efforts to train, attract, and retain the world's best AI minds. This includes scaling education and fellowship programs in AI, expediting visas for AI experts, and fostering public-private research alliances. AI talent concentration is a force multiplier for innovation; the U.S. must remain the **magnet for global AI expertise** to maintain leadership.

5. Prioritize Trustworthy AI and Alignment Research: Make AI safety an integral part of the R&D agenda. Dedicate funding to solve AI **factuality, robustness, and bias** issues – so our systems are reliable – and to develop alignment techniques that keep advanced AI behaviors in check. This ensures that progress in capability comes with commensurate progress in safety, sustaining public trust and ethical use of AI.

6. Maintain Strategic Leadership vis-à-vis China: Implement a coordinated national strategy to outpace China (and other competitors) in AI. This means **increasing federal AI R&D investment**, safeguarding chip supply chains, and leveraging alliances to amplify our advantages. We cannot let rivals gain a decisive edge in any critical AI domain – whether it's LLMs, military AI, or robotics. Constant vigilance and agility in policy will keep the U.S. ahead.

7. Leverage AI for Scientific and Technological Breakthroughs: Fully integrate AI into America's innovation enterprise. From AI-designed drugs and materials to autonomous research robots, use AI to accelerate discovery in health, energy, environment, and more. Launch a "National AI for Science" initiative that equips our scientists with powerful AI tools and establishes moonshot challenges (like an AI winning a Nobel) to drive progress.

8. Build Public-Private AI Partnerships: Strengthen collaboration between government, academia, and industry in AI development. This includes co-funded research projects (e.g. at national labs), shared testbeds, and open innovation challenges. The recent engagement of OpenAI with federal labs on advanced reasoning models shows the value of such partnerships. The federal government should act as convener and co-investor to unify the nation's AI efforts.

9. Modernize Research Funding and Infrastructure: Update the "plumbing" of innovation. Streamline grant processes to be more nimble (fast-track high-potential ideas, support interdisciplinary work) and ensure sustained funding for compute resources and data sharing. Appoint an **AI Infrastructure Czar** to

coordinate large projects and cut red tape. Modern methods and leadership will accelerate the pace from research to results.

10. Champion U.S. Values in AI Development: Finally, as we lead in AI R&D, ensure it reflects American values of openness, fairness, and human-centric progress. Establish ethical guidelines for AI, involve diverse stakeholders in setting research priorities, and promote global norms for responsible AI use. U.S. leadership is not just about being first – it's about setting the **standard** for AI that benefits all of society, safely and equitably.