HUGGING FACE

Hugging Face Response to Request for Information on the Development of an Artificial Intelligence Action Plan

March 2025

About Hugging Face

Hugging Face is a community-driven U.S. company that democratizes responsible artificial intelligence (AI) and machine learning (ML). Our platform is the most widely used for sharing and collaborating on ML systems, fostering open-source and open-science initiatives. We host machine learning models and datasets within an infrastructure that enables efficient data processing, analysis, and research. Additionally, we provide educational resources, courses, and tooling to lower the barrier to AI participation for individuals from all backgrounds.

Executive Summary

Hugging Face appreciates the opportunity to submit comments to the Office of Science and Technology Policy (OSTP) regarding the development of an Artificial Intelligence Action Plan. As a leading AI company committed to democratizing artificial intelligence through open-source and collaborative approaches, we believe that thoughtful policy can support innovation while ensuring that AI development remains competitive, and aligned with American values.

Based on Hugging Face's experience as a leading AI platform serving 7 million users and hosting over 1.5 million public models across diverse domains, we propose an AI Action Plan centered on three interconnected pillars:

1. **Strengthen Open and Open-Source AI Ecosystems**: technical innovation comes from diverse actors across institutions. Support for infrastructure like NAIRR and investment in open science and data allows these contributions to have an additive effect and accelerate robust innovation.
2. **Prioritize Efficient and Reliable Adoption**: spreading the benefits of the technology by facilitating its adoption along the value chain requires actors across sectors of activity to shape its development. More efficient, modular, and robust AI models require research and infrastructural investments to enable the broadest possible participation and innovation, enabling diffusion of technology across the U.S. economy.

3. **Promote Security and Standards**: Decades of practices in open-source software cybersecurity, information security, and standards can inform safer AI technology. Promoting traceability, disclosure, and interoperability standards will foster a more resilient and robust technology ecosystem.

Founded in 2016, Hugging Face has become the main platform supporting AI development through open-source contributions spanning language, biology, robotics, law, finance, and beyond. Our recommendations are grounded in practical experience developing tools that have become industry standards while supporting a diverse ecosystem of researchers, startups, and established companies.

The Role and Progress of Open and Transparent AI

Modern AI is built on decades of open research, with commercial giants relying heavily on open source contributions like PyTorch and groundbreaking research on transformer architectures, attention mechanisms, and training methodologies. Further, recent breakthroughs like OLMO2 – a relatively small model with fully transparent training methods and data matching OpenAI's o1-mini's performance – and OlympicCoder – an even smaller models exceeding the performance of the latest Claude model on complex coding problems – demonstrate that open research remains a promising path to developing systems that match the performance of commercial models, and can often surpass them especially in terms of efficiency and performance in specific domains. Perhaps most striking is the rapid compression of development timelines—what once required over 100B parameter models just two years ago can now be accomplished with 2B parameter models, suggesting an accelerating path to parity. This trend towards more accessible, efficient, and collaborative AI development shows that open approaches to AI development have a critical role to play in enabling a successful AI strategy that maintains technical leadership and supports more widespread and secure adoption of the technology. We go into further detail in the rest of this document.

Open Source and Open Science are Key to U.S. Leadership

Open models, infrastructure, and scientific practices constitute the foundation of AI innovation, allowing a diverse ecosystem of researchers, companies, and developers to build upon shared knowledge. Hugging Face's platform hosts AI models and datasets from both small actors (startups, universities) and large organizations (Microsoft, Google, OpenAI, Meta), demonstrating how open approaches accelerate progress and democratize access to AI capabilities. Our own research and software contributions, such as the Transformers and Diffusers libraries – which have become the standard framework for implementing and sharing

state-of-the-art models – exemplify how open science drives the field forward by enabling rapid iteration and validation across the global research community.

The United States must lead in open-source AI and open science, which can enhance American competitiveness by fostering a robust ecosystem of innovation and ensuring a healthy balance of competition and shared innovation. Research has shown that open technical systems act as force multipliers for economic impact, with an estimated 2000x multiplier effect — meaning $4 billion invested in open systems could potentially generate $8 trillion in value for companies using them. These economic benefits extend to national economies as well. Without any open-source software contributions, the average country would lose 2.2% of its GDP. Open source drove between €65 and 95 billion of European GDP in 2018 alone—a finding so significant that the European Commission cited it when establishing new rules to streamline the process for open-sourcing government software. This demonstrates how open-source impact translates directly into policy action and economic advantage at the national level, underlining the importance of open source as a public good.

The commercial adoption of open models is driven by several practical factors. First, cost efficiency — developing AI models from scratch requires significant investment, so leveraging open foundations reduces R&D expenses. Second, customization — organizations can adapt and deploy models specifically tailored to their use cases rather than relying on one-size-fits-all solutions. Third, reduced vendor lock-in—open models give companies greater control over their technology stack and independence from single providers. Finally, open models have caught up to and in certain cases, surpassed the capabilities of closed, proprietary systems: Olympic-Coder, released as part of Hugging Face's Open R1 project, surpasses Claude 3.7, the latest proprietary model from Anthropic, in terms of coding performance. All this is particularly valuable for startups and mid-sized companies, which can access cutting-edge technology without massive infrastructure investments. Banks, pharmaceutical companies and other industries have been adapting open models to specific market needs, demonstrating how open-source foundations support a vibrant commercial ecosystem across the value chain.

Policy Recommendations:

- **Enhance National Research Infrastructure**: Fully implement and expand the National AI Research Resource (NAIRR) pilot. Hugging Face's active participation in the NAIRR pilot has demonstrated the value of providing researchers with access to computing resources, datasets, and collaborative tools.
- **Allocate Public Computing Resources for Open Source**: The public should have ways to participate via public AI infrastructure. A way to do this would be to dedicate a portion of publicly-funded computing infrastructure to support open-source AI projects, reducing barriers to innovation for smaller research teams and companies that cannot afford proprietary systems.

- **Enable access to data for developing open systems:** Create sustainable data ecosystems through targeted policies that address the decreasing data commons. Publishers are increasingly signing data licensing deals with proprietary AI model developers to the tune of hundreds of millions of dollars, meaning that quality data acquisition costs are now approaching or even surpassing computational expenses of training frontier models, threatening to lock out small open developers from access to quality data. Support organizations that contribute to public data repositories and streamlined compliance pathways that reduce legal barriers to responsible data sharing.
    - **Develop Open, High-Quality Datasets**: Invest in the creation, curation, and maintenance of robust, representative datasets that can support the next generation of AI research and applications. Expand initiatives like the IBM AI Alliance Trusted Data Catalog and support projects like IDI's AI driven Digitization of the public collections in the Boston Public Library.
    - **Strengthen Rights-Respecting Data Access Frameworks**: Establish clear guidelines for data usage, including standardized protocols for anonymization, consent management, and usage tracking. Support public-private partnerships to create specialized data trusts for high-value domains like healthcare and climate science, ensuring that individuals and organizations maintain appropriate control over their data while enabling innovation.
- **Invest in Stakeholder-Driven Innovation**: Create and support programs that enable organizations across diverse sectors (healthcare, manufacturing, education) to develop customized AI systems for their specific needs, rather than relying exclusively on general-purpose systems from major providers. This enables broader participation in the AI ecosystem and ensures that the benefits of AI extend throughout the economy.
- **Strengthen Scientific Centers of Excellence**: Expand NIST's role as a convener for AI experts across academia, industry, and government to share lessons and develop best practices. In particular, the AI Risk Management Framework has played a significant role in identifying stages of AI development and research questions that are critical to ensuring more robust and secure technology deployment for all. The tools we develop at Hugging Face, from model documentation to evaluation libraries, are directly shaped by these questions.
- **Support High-Quality Data for Performance and Reliability Evaluation**: AI development depends heavily on data, both to train models and to reliably evaluate their progress, strengths, risks and limitations. Fostering greater access to public data in a safe and secure way and ensuring that the evaluation data used to characterize models is sound and evidence-based will accelerate progress in both performance and reliability of the technology.

HUGGING FACE

Prioritize Efficient and Reliable Adoption

Smaller companies and startups face significant barriers to AI adoption due to high costs and limited resources. According to IDC, global AI spending will reach $632 billion in 2028, but these costs remain prohibitive for many small organizations. Meanwhile, energy scarcity presents a growing concern, with the International Energy Agency projecting that data centers' electricity consumption could double from 2022 levels to 1,000 TWh by 2026 — equivalent to Japan's entire electricity demand.

While training AI models is energy-intensive, inference — due to its scale and frequency — can ultimately exceed training energy consumption. Ensuring broad AI accessibility requires both hardware optimizations and scalable software frameworks. A range of organizations are developing models tailored to their specific needs, and U.S. leadership in efficiency-focused AI development presents a strategic advantage. The DOE's AI for Energy initiative further supports research into energy-efficient AI, facilitating wider adoption without excessive computational demands.

Open-source AI tools also bring financial returns, 51% of surveyed companies currently utilizing open-source AI tools report seeing positive ROI, as compared to just 41% of those not using open source. Knowledge sharing and innovation enabling organizations to efficiently use their existing compute resources, such as Hugging Face's Ultra-Scale playbook, help organizations of all sizes efficiently train their own AI models on hardware that they already have access to without requiring massive infrastructure investments. Techniques like quantization, pruning, and model distillation significantly reduce computational requirements while maintaining performance, as demonstrated by Hugging Face's SmolLM and SmolVLM models. Lightweight frameworks further enhance AI deployment in edge devices and resource-constrained environments. Beyond language models, open-source robotics helps lower costs, supporting research and innovation in fields where hardware expenses present a barrier.

Reliability is equally important for widespread AI adoption. Organizations need systems that function consistently as intended, particularly in critical domains. Studies have shown error rates as high as 19% in some healthcare AI applications, while general-purpose AI systems often fail to meet the specific needs of specialized sectors, emphasizing the need for tailored solutions. Locally hosted open models are also, by design, more reliable - unlike proprietary models accessed via APIs, locally hosted models provide version stability that is often a requirement for critical use cases.

Policy Recommendations:

- **Invest in Efficient AI Research**: Support R&D investments toward developing models that maintain high performance while reducing computational and energy requirements for

inference and other non-training workflows, including methods to transfer large model capabilities to drastically smaller models and models that can run on edge devices.

- **Support Efficiency Measurements and Benchmarks**: Develop federal standards for measuring and reporting AI system efficiency, creating market incentives for technologies that use limited energy resources more effectively. Examples of tools include the Energy Star rating for models.
- **Encourage Appropriate Technology Selection**: Promote the principle of using the most efficient tool that meets performance requirements (such as resource utilization considerations in the NIST AI RMF), ensuring appropriate resource utilization across the AI ecosystem.
- **Support Cross Agency Collaboration of Expertise**: Create and continue to support mechanisms for AI expertise to be shared across government agencies, preventing siloed knowledge and enabling consistent approaches to AI policy and procurement.
- **Build Evaluation Capacity**: Provide training and further develop technical capabilities within organizations, including government agencies, to evaluate AI systems they procure or deploy, ensuring informed decision-making. This requires both technical infrastructure and human expertise to create, choose, and run evaluations, and to interpret evaluation results in context.
- **Support Representative Evaluation Infrastructure**: Invest in developing evaluation frameworks that accurately represent the variety of potential AI applications and users across different contexts. Developer-run private evaluations are not currently comparable and have limited scrutiny, whereas open evaluations enable broader oversight.

Promote Security and Standards

As AI systems become increasingly integrated in all aspects of digital infrastructure, ensuring that they meet appropriate standards of documentation and information security will be essential to ensuring the resilience of the entire ecosystem.

First, while AI systems do present some new challenges – particularly in terms of scale – they remain first and foremost digital information systems; so AI security practices should be informed by decades of experience in cybersecurity. These include for example ensuring sufficient documentation of all components of a system, so model cards and training information for AI systems can play a similar role for AI security as the widely recognized Software Bill Of Materials. As for software, recognizing the importance of open source AI as a necessary component of every critical infrastructure sector, and enabling Secure by Design development practices by ensuring federal agencies have capacity to work on every stage of the AI development chain will provide a strong foundation for a more secure digital ecosystem.

HUGGING FACE

Open science in AI and open source AI systems are particularly relevant to the long-term security of critical infrastructure in the U.S. Developing open infrastructure to democratize state-of-the-art AI training techniques enables organizations to train their own performant models in controlled environments. "Maximally open" AI models, which provide full transparency into their training datasets and processes, enable extensive investigation to support security certifications in critical settings. "Open-weight" models stored in secure formats that can be deployed in air-gapped environments play a critical role for adoption of the technology in the most sensitive environments.

Open and open source development also has a significant role to play in setting favorable standards. One motivation for developing open source software for both private and government actors has been its ability to shape global practices. Work from international partners such as the UK AI Security Institute or Singapore's Digital Trust Center on open tooling has proven influential. Investing in open source software will support U.S. leadership in global standard-setting processes. In light of these considerations, prioritizing open and open-source AI as a critical component of the U.S. AI security strategy will be crucial for its success.

Policy Recommendations:

- **Support Strong Documentation and Transparency Standards:** Ensure that all actors using AI systems in security-impacting applications have sufficient information about the system's development and characteristics to proactively identify risks in their deployment context.
- **Build Internal Capacity to Develop Secure AI Systems:** Facilitate the development of purpose-specific AI systems for government use in cases where the use of general-purpose commercial systems may carry information security risks.
- **Develop Public-Private Evaluation Partnerships that Prioritize Open Development:** Establish frameworks for government agencies to collaborate with private-sector AI developers on evaluation efforts, combining government use cases with industry technical expertise, with an explicit mandate to prioritize open evaluation data and tooling – building on initiatives like the AI Safety Institute Consortium.
- **Extend Existing Cybersecurity Practices to AI Systems:** Ensure that cybersecurity and resilience of the software and data transfers underlying AI development and deployment remain a priority, in particular by leveraging decades of experience in securing open source systems.

HUGGING FACE

Conclusion

Continuing to lead AI research, development, and deployment in a direction that promotes innovation, efficiency, and competition requires that we act now. Open technical systems, from open science and open source AI, to open software tools, boost economies and accelerate innovation, providing a strong incentive for the U.S. to maintain leadership in open AI development. This would strengthen the growing ecosystem of AI applications that meet the needs of the entire range of actors in the U.S. economy by balancing the complementary relationship between open and proprietary innovation.

By leveraging the strengths of open-source approaches, prioritizing efficiency, and promoting interoperability standards, we are confident that the AI Action Plan will foster an environment where American companies and researchers thrive while developing technologies that benefit society broadly.

Hugging Face remains committed to contributing to this ecosystem through our open-source tools, collaborative platforms, and research initiatives. We appreciate the opportunity to provide these recommendations and look forward to continuing engagement with the OSTP and government stakeholders in shaping effective AI policy.

Submitted by:
Avijit Ghosh, Yacine Jernite, and Irene Solaiman
Hugging Face