

# Open Weights and American AI Leadership

July 24, 2026

In the 1980s, early open-source software pioneers challenged the prevailing belief that software would advance only if companies kept tight control over their code. This movement pushed for a transparent ecosystem where developers around the world could study, modify, and improve software. Software developed by the open-source community now supports most of the internet and underlies systems used by the world's largest technology companies, as well as the U.S. military and federal agencies conducting scientific research, cybersecurity, and other critical missions. Open source did more than lower the cost of software; it created a shared foundation of knowledge on which generations of American engineers and entrepreneurs built their institutional sovereignty.

The United States now faces a similar choice with artificial intelligence. Our AI leadership will be judged not by one frontier AI model, but by whether the United States builds a strong, open ecosystem that diffuses into every sector. This is essential for creating opportunities for innovation and prosperity across the country. It requires expanding access to AI, encouraging competition, robust application layers, and giving Americans greater control over the technology they rely on. Open weight models—AI models that anyone can download, inspect, modify, and run on their own infrastructure—are an important part of that foundation because they make advanced AI more accessible, adaptable, and widely available.

Open weights expand access to the AI economy. Startups, established businesses, universities, and public institutions can build on advanced models without training one from scratch or paying frontier-model prices for every task. Open weights let every organization match the right model to the right job at the right cost, reserving frontier-scale capability for genuine frontier problems and running efficient, specialized models everywhere else. That discipline is what will make AI economically sustainable as its use scales into the billions of everyday tasks. America wins the AI era by diffusing it into the

workflows of factories, hospitals, farms, classrooms, and main street businesses.

Open weights also strengthen competition and competition is what keeps the gains of AI broadly shared rather than concentrated in a few hands. By allowing many organizations to build, adapt, and deploy advanced models, open weights create rivalry not only among model developers but across cloud chips, applications, and services. That competition spurs innovation, drives down costs, and distributes the benefits of AI broadly across our economy.

Open weights also give customers greater control. As organizations invest in AI, they want to know that they will not become locked into a single provider or lose the knowledge and capabilities they build over time. Open weight models help provide that assurance by allowing organizations to control their own data, evaluate and adapt models to their own needs, and deploy them wherever their business requirements demand. And as organizations create value with AI, open weights allow them to own that value through self-improving models, specialized capabilities, and accumulated knowledge that drive American sovereignty and prosperity.

To be sure, open weights carry real and distinct risks. Once released, the weights are beyond the original developer's control, and modified versions are difficult to trace or reverse. But the right response to this risk is not to prohibit open weights. In a world where cybersecurity attackers use advanced AI, defenders need access to models with comparable capabilities so they can detect, simulate, and respond to emerging threats. Open models broaden defensive capability, increase transparency, and allow vulnerabilities to be discovered and remediated across many teams.

In fact, openness may be one of the most important paths to AI safety and security. Relying solely on closed models is not inherently safe: they can be breached, misused, or fail in ways that outsiders cannot detect. And concentrating advanced AI capabilities behind a small number of closed models compounds that risk. It results in a small number of single points of failure, weakens competition, and leaves critical technology in the hands of a few providers. Open weight models, on the other hand, allow a broad community of researchers and developers to examine their behavior, identify vulnerabilities, develop safeguards, and improve them over time. Just as open-source software demonstrated that transparency can be more secure than obscurity, AI safety may depend on giving more people the ability to test and strengthen the models on which society relies. It allows for rigorous benchmarking and evaluation, red teaming, and protections tied to real and demonstrated harms rather than assuming that closed systems are safer by default.


A strong AI ecosystem is not a foregone conclusion. Policymakers have an important opportunity to act. This includes expanding access to compute for startups and researchers, investing in shared training assets (datasets, tools, evaluation frameworks), and keeping the frontier plural by avoiding premature restrictions on open models that stifle competition or drive innovation overseas. These measures must also look at how strong application layers can expand sovereign use of AI across the economy.


In shaping this ecosystem, policymakers should be careful not to conflate legitimate model-development techniques with misappropriation. Distillation, or the practice of using one model's outputs to help train or improve another, is a widely used technique for model improvement, evaluation, and validation. It reflects a long tradition of learning from, building upon, and improving existing technologies, a tradition that has helped drive innovation since the rise of the open-source software movement. By contrast, unlawful efforts to extract value from closed models raise legitimate concerns. Those concerns should be addressed through targeted legal and commercial frameworks rather than sweeping restrictions on techniques that play an important role in AI innovation.

The age of AI can be one of prosperity. With the right choices, open weight AI can expand opportunity, strengthen competition, extend American technological leadership, mitigate risk, and ensure that the benefits of this extraordinary technology are shared broadly across our economy. That future is worth building, and the United States should lead in building it.

## Signatories

Agno • AI21 • alphaXiv • AMD • American Innovators Network • Amazon • AMP •  
 Andreessen Horowitz • AnythingLLM • Applied Compute • Arcee AI • Arena • ARK  
 Invest • Armada • Atreides Management • Automattic • Baseten • Black Forest Labs •  
 Block • Bolt • Box • Bria • Browserbase • Camber • Cisco • Cline • Cloudflare • Cohere •  
 Comcast • Core Automation • CoreWeave • Corridor • CrowdStrike • Crusoe •  
 Databricks • DataRobot • DatologyAI • Daytona • Dell Technologies • Digital Ocean •  
 Disruptive • DoorDash • E2B • EdgeRunner • EleutherAI • Emergence Capital •  
 Emergent • Exia Labs • Fastino Labs • Fireworks AI • FriendliAI • Genspark • GitHub •  
 GitLab • Glean • GMI Cloud • Goodfire • Google • GPU MODE • Groq • Harness •  
 Hugging Face • humans& • IBM • Inferact • Intangible • Intel • Interconnects AI •  
 LangChain • Letta • Lightning AI • Liquid AI • LM Studio • LMSYS • Mariana Minerals •  
 Merge • Meta • Microsoft • Mistral • Modal • Morph • Mozilla • Nebius • Neon • Netflix •  
 Notion • Nous Research • NVIDIA • Ollama • OpenAI • OpenClaw • OpenCode •  
 Palantir • Palo Alto Networks • Periodic Labs • Perplexity • Plastic Labs • Postman •  
 Prime Intellect • PrismML • RadixArk • Red Hat • Reducto • Reflection • Rehearsals •  
 Reka • Replit • Resemble AI • Recursive Intelligence • Runway • Sakana AI • SAP • Scale •  
 ServiceNow • Siemens • SkyPilot • SpaceX • Stack Overflow • Sycamore • Telnyx •  
 TensorWave • The Linux Foundation • Together AI • TrainLoop • Trajectory • TrendAI •  
 Uber • Unsloth • Unusual Ventures • Vercel • WRITER • Y Combinator • Zoom





scale

servicenow

SIEMENS

SkyPilot

SPACEX

Stack Overflow

Sycamore

telnyx

TENSORWAVE

together.ai

trainloop

Trajectory

TrendAI

Uber

unsloth

UNUSUAL

Vercel

WRITER

Y

zoom

## What's new

Surface Pro

Surface Laptop

Surface Laptop Ultra

Surface RTX Spark Dev Box

Copilot for organizations

Copilot for personal use

Explore Microsoft products

Windows 11 apps

## Microsoft Store

Account profile

Download Center

Microsoft Store support

Returns

Order tracking

Certified Refurbished

Microsoft Store Promise

Flexible Payments

## Education

Microsoft in education

Devices for education

Microsoft Teams for Education

Microsoft 365 Education

How to buy for your school

Educator training and development

Deals for students and parents

AI for education

## Business

Microsoft AI

Microsoft Security

Dynamics 365

Microsoft 365

Microsoft Power Platform

## Developer & IT

Azure

Microsoft Developer

Microsoft Learn

Support for AI marketplace apps

Microsoft Tech Community

## Company

Careers

About Microsoft

Company news

Privacy at Microsoft

Investors

|                       |                       |                         |
|-----------------------|-----------------------|-------------------------|
| Microsoft Teams       | Microsoft Marketplace | Diversity and inclusion |
| Microsoft 365 Copilot | Software companies    | Accessibility           |
| Small Business        | Visual Studio         | Sustainability          |

 English (United States)

 Your Privacy Choices

Consumer Health Privacy

[Sitemap](#)   [Contact Microsoft](#)   [Privacy](#)   [Manage cookies](#)   [Terms of use](#)   [Trademarks](#)   [Safety & eco](#)   [Recycling](#)  
[About our ads](#)   © Microsoft 2026