

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.

1789 Capital • A.Team • Adaption • Agno • AI Native Foundation • AI Tinkerers • AI21 • alphaXiv • Amazon • AMD • American Innovators Network • AMP • Amplitude • Anaconda • Andreessen Horowitz • AnythingLLM • Applied Compute • Arcee AI • Arena • ARK Invest • Armada • Assembled • Atlassian • Atomic Chat • Atreides Management • Autohand AI • Automattic • Baseten • Black Forest Labs • BLACKBOX AI • Block • Bolt • Box • Bria • Browserbase • Cadence • Camber • Capy • Cinder • Ciroos • Cisco • Cline • CloudBees • Cloudflare • Cogensec • Cognition • Cohere • Cohesity • Coinbase • Comcast • Command Code • Composio • Core Automation • CoreWeave • Corridor • Corvic AI • Criteria Corp • CrowdStrike • Crusoe • DART • Databricks • DataRobot • DatologyAI • Daytona • Defense Unicorns • Dell Technologies • depthfirst • DigitalOcean • Disruptive • DoorDash • E2B • EdgeRunner • Elastic • EleutherAI • Elorian AI • Emergence Capital • Emergent • Endex • EOX Vantage • Exia Labs • Factory • Fastino Labs • Featherless AI • Fireworks AI • Fleet AI • FriendlyAI • General Catalyst • Genspark • GitHub • GitLab • Glean • Glemad • GMI Cloud • GoDaddy • Goodfire • Google • GPU MODE • Groq • Harness • Hugging Face • humans& • Hyde • Hydra Host • Hyland • Hyperbolic • IBM • Inferact • InferX • Intangible • Intel • Interconnects AI • Isomer • Kindo • Kong • Kyndryl • Lambda • LangChain • Latitude.sh • Lenovo • Letta • Lightning AI • Liquid AI • LithosAI • LM Studio • LMSYS • Makora • Mariana Minerals • Megaport • Megaton AI • Merge • Meta • microagi • Microsoft • Mistral • Modal • Morph • Mozilla • Nebius • Neon • NeoSigma • Netlify • Noah Labs • Nokia • Nolla Health • Notion • Nous Research • NVIDIA • Ollama • Open Athena • Open Compute Project Foundation • Open WebUI • OpenAI • OpenClaw • OpenCode • Osmantic • Oumi • Palantir • Palo Alto Networks • Periodic Labs • Perplexity • Pinterest • Plastic Labs • PlayerZero • Poolside • Postman • Prime Intellect • PrismML • Rackspace Technology • RadixArk • Red Hat • Redblock • Redis • Reducto • Reflection • Regal • Rehearsals • Reka • Replit • Resemble AI • Recursive Intelligence • RightNow AI • Rillet • Rogo • Runway • Sakana AI • San Francisco Tooling Company • SAP • Sapiom • Scale • Scarf • Sculptor AI • Seekr • SerpApi • ServiceNow • SF Tensor • Siemens • SkyPilot • Snowflake • Socket • Sonar • SpaceX • Stack Overflow • Stacksync • Sully.ai • Superhuman • Supermemory • Superserve • Sycamore • Synopsys • Telnix • TensorWave • The Linux Foundation • Together AI • Tolmo • TrainLoop • Trajectory • TrendAI • Uber • Uniphore • Unsloth • Unusual Ventures • Vercel • Verda • Vocara • Webflow • Willow • WordPress.org • Workday • WRITER • Wrynx • xpander.ai • Y Combinator • Zendesk • Zoom • Zscaler



