

FREE EXPRESSION

A Daily Newsletter From WSJ Opinion

ABOUT FREE EXPRESSION

SIGN-UP FOR THE FREE NEWSLETTER

CONTINETTI: TRUMP SH

This copy is for your personal, non-commercial use only. Distribution and use of this material are governed by our Subscriber Agreement and by copyright law. For non-personal use or to order multiple copies, please contact Dow Jones Reprints at 1-800-843-0008 or visit www.djreprints.com.

<https://www.wsj.com/opinion/america-needs-to-go-on-offense-with-ai-54744c02>

OPINION COMMENTARY [Follow](#)

America Needs to Go on Offense With AI

Grad students love open Chinese models. The U.S. should be competing, not seeking to ban them.

By Jonathan Zhukovsky

Aug. 6, 2026 2:18 pm ET



CHEN HAOMING/XINHUA/ZUMA PRESS

Nvidia CEO Jensen Huang joined X in June, and his first post, two weeks ago, wasn't a product announcement. It was a policy letter. Twenty-five companies, including Nvidia, Microsoft and Palantir, urged Washington not to restrict open-weight artificial-intelligence models—those whose underlying parameters anyone can download, inspect and adapt. When the chief executive of the world's most valuable company (or second-most, depending on the day) uses his social-media debut to publish a plea to policymakers, it's worth asking what he's worried about.

The immediate answer is a reported push inside the Trump administration to ban Chinese open models. The proximate cause is Moonshot AI's Kimi K3, released last month: 2.8 trillion parameters, the largest open model ever published, benchmarking within reach of the best American proprietary systems. Global demand overwhelmed Moonshot's computing capacity within 48 hours.

But the would-be banners and the letter's signatories are arguing about the wrong thing. They treat AI models as files to be restricted or protected, ignoring the only scoreboard that matters: which models the grad students at Berkeley, Stanford, MIT and Carnegie Mellon reach for when they start an experiment.

Right now, the answer should alarm anyone who cares about American technological leadership. According to the ATOM Report, published this spring by American Truly Open Models, Chinese models overtook American ones in cumulative downloads in the summer of 2025 and have since opened a gap of more than 400 million. By early this year, 70% of all new fine-tuned models—the derivative works researchers and developers build on top of base models—were built on Chinese foundations. Alibaba's Qwen family alone accounts for 69% of new fine-tunes, while Meta's Llama has collapsed from 44% to 11%. On inference platforms, Chinese models' share of tokens served rose from under 3% to over 70% in 14 months.

This isn't because our researchers are disloyal. Chinese labs release new and improved models every few months, so researchers can count on the family they've invested in staying current—while American open releases have been sporadic, often arriving once and going quiet. Chinese open-weight models run on everything from a student's laptop to a supercomputer, with no legal strings attached. And within days of each release, a global community of contributors builds the supporting software—the add-ons and utilities that make a model practical to use. American AI research is being built, as a result, on Chinese foundations.

In the 1990s, the software establishment believed serious systems required tight proprietary control. Today open-source code underpins the internet,

the largest technology companies and the Pentagon's own systems. The lesson wasn't that openness beats secrecy. The winner of a platform war is whoever captures the contribution loop—the self-reinforcing cycle in which researchers, toolmakers and downstream builders improve a platform. Linux, a popular open-source operating system, didn't win on license terms. It won in the community.

China's AI labs have internalized this lesson with remarkable discipline. DeepSeek, Alibaba, Zhipu and Moonshot release open-weight models not out of ideological commitment but as an industrial strategy: Every American graduate student who fine-tunes Qwen and every startup that builds on Kimi extends Beijing's technological gravity at zero marginal cost. Open-sourcing is how a nation exports its stack.

America's response so far has been defensive. Ban the Chinese models. Restrict the downloads. But you don't win a platform war on defense, and prohibition here is worse than useless. A ban won't make a Berkeley doctoral student reach for an American model that doesn't meet his needs; it will impede his research, force him to work offshore or go underground, and confirm to every allied government watching that American tech comes with political strings attached.

If allied nations and American enterprises conclude that the vibrant, dependable open ecosystem is the Chinese one, they will build their sovereign AI infrastructure on it. I evaluate open models weekly for exactly such deployments—isolated, sovereign systems for which the community behind a model matters as much as its benchmarks—and I can report that the gap between American press releases and Chinese ecosystems is real.

What is required for the U.S. to win? Five things, none of which is a ban.

First, top-tier American models must be released openly and regularly—not as one-time gestures. A community can't form around a model that shows up once and is abandoned. Second, labs should release each model in multiple open sizes, as Chinese labs already do, so what is learned on small models carries over as researchers scale up. Third, the raw, unpolished

versions of models must be released—not just the finished consumer products. Researchers need the raw material to build on. Fourth, we need licenses that let people freely use and modify what they build. Meta’s Llama models attracted a real following, then lost it with legal fine print that scared off lawyers and businesses. Fifth, government and industry alike should fund the shared plumbing of an open-weight environment—testing tools, public datasets, computing power for universities—so that building on American models is the easy choice, not an act of patriotism.

Notably absent from Mr. Huang’s letter were OpenAI and Anthropic—the labs most capable of releasing frontier-class American open models, but also among those lobbying for restrictions on Chinese ones. There is something suspicious about our leading labs keeping their best work closed while asking Washington to handicap the competition’s open work. That isn’t how we won the software era.

When American labs ship serious open models, the community responds: Nvidia’s Nemotron is adopted at 17 to 20 times the pace of the rest of its size class, according to the ATOM report. The demand for American open models is enormous. The supply isn’t.

The day the typical AI paper out of Stanford fine-tunes an American open model—simply because it is the best model available—is the day we are winning. Until then, every export-control hearing in Washington is a sideshow. The real scoreboard is in Berkeley, and we are behind.

Mr. Zhukovsky is chief technology officer of Hillspire, family office of Eric and Wendy Schmidt.

Appeared in the August 7, 2026, print edition as ‘America Needs to Go on Offense With AI’.

Further Reading

Inside Amazon’s AI Security Playbook

Wall Street Thinks It Knows How Tech Giants Will Make AI Pay

Amazon Cuts Jobs in Artificial General Intelligence Unit

The Quarter-Trillion-Dollar Onslaught of AI Bonds Is Testing Investors' Limits

OPINION Will Washington State Choose Stagnation Over Growth?

Videos