FORTUNE®

JULY 3, 2026

COMMENTARY / SOFTWARE

I argued with the father of open source for 2 years. Now the AI fight is the same — only bigger.

By David Siegel

In the 1980s, I had the chance to spend several years arguing about free and open software, what we now call open source, with the founder of the movement, Richard Stallman. My office at the MIT AI Lab was next door to his. Stallman's position was that the source code to software should be free for everyone to use, learn from, and improve. Software encapsulates knowledge, he argued, and no one should lock something so fundamental away. To hide software inside a company was to hide knowledge itself.

At first I took the conventional view of the time. Software would only advance, I insisted, if companies kept proprietary control over their code. We agreed on the bigger picture—that computers would become a central accelerator of human progress—but disagreed sharply on how to get there. What I missed was that software was not just a commercial asset; it was a body of knowledge, and bodies of knowledge grow stronger when they are shared. After about two years of on-and-off debate, Stallman convinced me I was wrong.

Over the following years, Stallman turned that conviction into a movement. He argued that people should be able to study the software they used, change it, improve it, and share it with others. That became the free software movement, and it later helped form the foundation for what the world now calls open source.

He wrote GCC, the translation program that to this day takes the world's code and turns it into the machine language computers actually run. Its success came not only from Stallman's own work but from the thousands who contributed to it. The same spirit produced GNU/Linux, the operating system that now powers most of the internet. The modern world runs on open source, powered by the principles of open development.

Plenty of arguments were made at the time for why all of this was a terrible idea. At the top of the list was security: to keep computers safe, the thinking went, you had to hide their software. There is a certain logic to "security through obscurity," but the counterargument turned out to be stronger: transparency lets a worldwide community of developers find and fix problems, while obscurity only hopes no one looks hard. Openness won that argument decisively.

This openness also accelerated the tech industry in a less obvious way. A vibrant open source community spread the knowledge of how to build, serving as the definitive textbook from which a generation of engineers learned. Conversely, a locked-down system teaches almost no one. How can we train the next generation of innovators if every state-of-the-art system is hidden from view?

The decades since have produced an extraordinary run of technological progress, powered by a delicate ecosystem—private companies, universities and tens of thousands of volunteers all contributing to a shared base of open software. Some of the world's most successful companies are built on it, earning their profits on services layered atop an open core. Proprietary software has its place. My point is narrower and harder to dispute: open source has been load-bearing, and we should be careful before we let it erode.

Because that is what is now happening. AI is software, and AI is increasingly closed. The frontier models—the most advanced, cutting-edge AI systems—are closed completely and the trend is accelerating. Viable open alternatives are few and far between. And this is happening while the science is still young: today's models are remarkable but unfinished, and the methods behind them are far from settled. Closing the field down now, with the deepest breakthroughs still ahead, is exactly when shared knowledge matters most.

This matters far beyond the software industry. Stallman's instincts were shaped by university science, where research is published openly for the world to build on. The advancement of society depends on exactly that kind of sharing. If most future science comes to rely on AI, then locking AI inside a few companies risks locking down much of scientific progress along with it.

Think about what a library is: a commons where anyone can walk in and draw on the whole inheritance of human knowledge, free of charge. We would find it intolerable if a few companies bought up every library, decided which books we were allowed to read, and quietly rewrote the ones we could. Yet AI is fast becoming the library of the future—and a closed, controlled AI is exactly that: a library you may enter only on the owner's terms.

There is a deeper problem hiding inside that one. When a handful of companies decide what a model can and cannot do, or quietly shape how it reaches its answers, then anyone who depends on it can no longer fully understand the results it produces. The same trap awaits the doctor leaning on a model for a diagnosis, the engineer trusting it with a design, the judge consulting it on a decision, and the ordinary person asking it what to believe.

One might object that this hardly matters, since the model can simply explain itself. But an explanation is not an audit. A model's stated reasons are a plausible story assembled after the fact, not a faithful record of the computation that produced the answer. Ask it the same question next year and you may get a different answer, with no way to know whether the world changed or the vendor did. Whoever depends on such a system isn't reasoning from a tool they comprehend; they are trusting an oracle they are not allowed to examine.

A standard objection is that AI is simply too dangerous to leave in the open. It's an argument worth taking seriously, because it isn't entirely crazy—it deserved a real answer. Critics argue that releasing the underlying AI is not like publishing a research paper: a paper describes a capability, while the software itself is the capability. That asymmetry is real.

But the conclusion doesn't follow. We could say the same of science itself—who knows what a published result might enable in the wrong hands?—yet we don't respond by classifying physics. We monitor, we set rules, and we keep the foundation open. Closed models are not safe by virtue of being closed; they leak, they get jailbroken, and their concentration creates its own danger—a few firms deciding what the rest of us are permitted to build. The honest question is not whether open models carry risk but whether they add meaningful risk beyond what is already available.

A key distinction is worth making carefully. There are two kinds of code behind any model: the code that runs it, and the code that built it. Being able to run a model is genuinely useful. But for almost everything I have argued here, what matters most is being able to see how it was built: the code that made it and the data it learned from. That transparency is vanishing fastest.

The models celebrated as open today—from China's leading labs and some American companies—do release the code to run them. What they hold back is the code that built them and the data their models learned from. So what you really get is a vast pile of numbers that somehow produces intelligence, with little account of how it came to be: magic numbers you can run but cannot explain. And even this much is a courtesy, not a commitment: the companies releasing these models make no promise to keep doing so for their most capable systems tomorrow.

An openness that can be switched off at will is not a foundation; it is a favor. We should want both the open models themselves, so anyone can use and build on these systems, and the open source, so anyone can see how they were made.

None of this is an argument against the companies building AI. I'm glad they exist. They simply should not be the only option. Private AI in America needs no defending while open source AI has no such champion and is the easiest thing to forget.

The path forward is not complicated, though it isn't free. Yes, frontier models keep getting bigger and more expensive—that arms race may well stay with the giants. But open source AI does not have to match their scale to be useful. Much of what the world needs probably does not require the absolute frontier. And where keeping a credible open option does demand serious compute, that is precisely the kind of public good worth paying for. What's missing is not a path but will. The government, the private sector, and nonprofits should invest heavily in free and open source AI—the way they once invested in open software: public compute grants for open research, corporate and philanthropic support for universities and nonprofits doing the work, and a simple rule that AI built with public money is open by default. We have run this experiment before. We know how it turns out. Let's not unlearn it.

The opinions expressed in Fortune.com commentary pieces are solely the views of their authors and do not necessarily reflect the opinions and beliefs of Fortune.

David Siegel is a computer scientist, entrepreneur, and philanthropist. He is co-founder of Two Sigma, founder of Open Athena, and chairman of Siegel Family Endowment.