



Cory Doctorow: What Kind of Bubble is AI?

📅 December 18, 2023 👤 Cory Doctorow



Cory Doctorow (by Amelia Beamer)

Of course AI is a bubble. It has all the hallmarks of a classic tech bubble. Pick up a rental car at SFO and drive in either direction on the 101 – north to San Francisco, south to Palo Alto – and every single billboard is advertising some kind of AI company. Every business plan has the word “AI” in it, even if the business itself has no AI in it. Even as two major, terrifying wars rage around the world, every newspaper has an above-the-fold AI headline and half the stories on Google News as I write this are about AI. I’ve had to make rule for my events: The first person to mention AI owes everyone else a drink.

It’s a bubble.

Tech bubbles come in two varieties: The ones that leave something behind, and the ones that leave *nothing* behind. Sometimes, it can be hard to guess what kind of bubble you’re living through until it pops and you find out the hard way.

When the dotcom bubble burst, it left a *lot* behind. Walking through San Francisco’s Mission District one day in 2001, I happened upon a startup founder who was standing on the sidewalk, selling off a fleet of factory-wrapped Steelcase Leap chairs (\$50 each!) and a dozen racks of servers with as much of his customers’ data as I wanted (\$250 per server or \$1000 for a rack). Companies that were locked into sky-high commercial leases scrambled to

sublet their spaces at bargain-basement prices. Craigslist was glutted with foosball tables and Razor scooters, and failed dotcom T-shirts were up for the taking, by the crateful.

But the most important residue after the bubble popped was the millions of young people who'd been lured into dropping out of university in order to take dotcom jobs where they got all-expenses paid crash courses in HTML, Perl, and Python. This army of technologists was unique in that they were drawn from all sorts of backgrounds – art-school dropouts, humanities dropouts, dropouts from earth science and bioscience programs and other disciplines that had historically been consumers of technology, not producers of it.

This created a weird and often wonderful dynamic in the Bay Area, a brief respite between the go-go days of Bubble 1.0 and Bubble 2.0, a time when the cost of living plummeted in the Bay Area, as did the cost of office space, as did the cost of servers. People started making technology because it served a need, or because it delighted them, or both. Technologists briefly operated without the goad of VCs' growth-at-all-costs spurs.

The bubble was terrible. VCs and scammers scooped up billions from pension funds and other institutional investors and wasted it on obviously doomed startups. But after all that "irrational exuberance" burned away, the ashes proved a fertile ground for new growth.

Contrast that bubble with, say, cryptocurrency/NFTs, or the complex financial derivatives that led up to the 2008 financial crisis. These crises left behind very little reusable residue. The expensively retrained physicists whom the finance sector taught to generate wildly defective risk-hedging algorithms were not able to apply that knowledge to create successor algorithms that were useful. The fraud of the cryptocurrency bubble was far more pervasive than the fraud in the dotcom bubble, so much so that without the fraud, there's almost nothing left. A few programmers were trained in Rust, a very secure programming language that is broadly applicable elsewhere. But otherwise, the residue from crypto is a lot of bad digital art and worse Austrian economics.

AI is a bubble, and it's full of fraud, but that doesn't automatically mean there'll be nothing of value left behind when the bubble bursts. WorldCom was a gigantic fraud and it kicked off a fiber-optic bubble, but when WorldCom cratered, it left behind a lot of fiber that's either in use today or waiting to be lit up. On balance, the world would have been better off without the WorldCom fraud, but at least something could be salvaged from the wreckage.

That's unlike, say, the Enron scam or the Uber scam, both of which left the world worse off than they found it in every way. Uber burned \$31 billion in investor cash, mostly from the Saudi royal family, to create the illusion of a viable business. Not only did that fraud end up screwing over the retail

investors who made the Saudis and the other early investors a pile of money after the company's IPO – but it also destroyed the legitimate taxi business and convinced cities all over the world to starve their transit systems of investment because Uber seemed so much cheaper. Uber continues to hemorrhage money, resorting to cheap accounting tricks to make it seem like they're finally turning it around, even as they double the price of rides and halve driver pay (and still lose money on every ride). The market can remain irrational longer than any of us can stay solvent, but when Uber runs out of suckers, it will go the way of other pump-and-dumps like WeWork.

What kind of bubble is AI?

Like Uber, the massive investor subsidies for AI have produced a sugar high of temporarily satisfied users. Fooling around feeding prompts to an image generator or a large language model can be fun, and playful communities have sprung up around these subsidized, free-to-use tools (less savory communities have also come together to produce nonconsensual pornography, fraud materials, and hoaxes).

The largest of these models are *incredibly* expensive. They're expensive to make, with billions spent acquiring training data, labelling it, and running it through massive computing arrays to turn it into models.

Even more important, these models are expensive to run. Even if a bankrupt AI company's model and servers could be acquired for pennies on the dollar, even if the new owners could be shorn of any overhanging legal liability from looming copyright cases, even if the eye-watering salaries commanded by AI engineers collapsed, the electricity bill for each query – to power the servers and their chillers – would still make running these giant models *very* expensive.

Do the potential paying customers for these large models add up to enough money to keep the servers on? That's the 13 trillion dollar question, and the answer is the difference between WorldCom and Enron, or dotcoms and cryptocurrency.

Though I don't have a certain answer to this question, I am skeptical. AI decision support is potentially valuable to practitioners. Accountants might value an AI tool's ability to draft a tax return. Radiologists might value the AI's guess about whether an X-ray suggests a cancerous mass. But with AIs' tendency to "hallucinate" and confabulate, there's an increasing recognition that these AI judgments require a "human in the loop" to carefully review their judgments.

In other words, an AI-supported radiologist should spend exactly the same amount of time considering your X-ray, and then see if the AI agrees with their judgment, and, if not, they should take a closer look. AI should make radiology

more expensive, in order to make it more accurate.

But that's not the AI business model. AI pitchmen are explicit on this score: The purpose of AI, the source of its value, is its capacity to increase productivity, which is to say, it should allow workers to do *more*, which will allow their bosses to fire some of them, or get each one to do more work in the same time, or both. The entire investor case for AI is "companies will buy our products so they can do more with less." It's *not* "business customers will buy our products so their products will cost more to make, but will be of higher quality."

AI companies are implicitly betting that their customers will buy AI for highly consequential automation, fire workers, and cause physical, mental and economic harm to their own customers as a result, somehow escaping liability for these harms. Early indicators are that this bet won't pay off. Cruise, the "self-driving car" startup that was just forced to pull its cars off the streets of San Francisco, pays 1.5 staffers to supervise every car on the road. In other words, their AI replaces a single low-waged driver with 1.5 more expensive remote supervisors – and their cars still kill people.

If Cruise is a bellwether for the future of the AI regulatory environment, then the pool of AI applications shrinks to a puddle. There just aren't that many customers for a product that makes their own high-stakes projects better, but more expensive. There are many low-stakes applications – say, selling kids access to a cheap subscription that generates pictures of their RPG characters in action – but they don't pay much. The universe of low-stakes, high-dollar applications for AI is so small that I can't think of anything that belongs in it.

Add up all the money that users with low-stakes/fault-tolerant applications are willing to pay; combine it with all the money that risk-tolerant, high-stakes users are willing to spend; add in all the money that high-stakes users who are willing to make their products *more* expensive in order to keep them running are willing to spend. If that all sums up to less than it takes to keep the servers running, to acquire, clean and label new data, and to process it into new models, then that's it for the commercial Big AI sector.

Just take one step back and look at the hype through this lens. All the big, exciting uses for AI are either low-dollar (helping kids cheat on their homework, generating stock art for bottom-feeding publications) or high-stakes and fault-intolerant (self-driving cars, radiology, hiring, etc.).

Every bubble pops eventually. When this one goes, what will be left behind?

Well, there will be little models – Hugging Face, Llama, etc – that run on commodity hardware. The people who are learning to "prompt engineer" these "toy models" have gotten far more out of them than even their makers

imagined possible. They will continue to eke out new marginal gains from these little models, possibly enough to satisfy most of those low-stakes, low-dollar applications. But these little models were spun out of big models, and without stupid bubble money and/or a viable business case, those big models won't survive the bubble and be available to make more capable little models.

There are some promising avenues, like “federated learning,” that hypothetically combine a lot of commodity consumer hardware to replicate some of the features of those big, capital-intensive models from the bubble's beneficiaries. It may be that – as with the interregnum after the dotcom bust – AI practitioners will use their all-expenses-paid education in PyTorch and TensorFlow (AI's answer to Perl and Python) to push the limits on federated learning and small-scale AI models to new places, driven by playfulness, scientific curiosity, and a desire to solve real problems.

There will also be a lot more people who understand statistical analysis at scale and how to wrangle large amounts of data. There will be a lot of people who know PyTorch and TensorFlow, too – both of these are “open source” projects, but are effectively controlled by Meta and Google, respectively. Perhaps they'll be wrestled away from their corporate owners, forked and made more broadly applicable, after those corporate behemoths move on from their money-losing Big AI bets.

Our policymakers are putting a lot of energy into thinking about what they'll do if the AI bubble *doesn't* pop – wrangling about “AI ethics” and “AI safety.” But – as with all the previous tech bubbles – very few people are talking about what we'll be able to salvage when the bubble is over.

*Cory Doctorow is the author of **Walkaway**, **Little Brother**, and **Information Doesn't Want to Be Free** (among many others); he is the co-owner of Boing Boing, a special consultant to the Electronic Frontier Foundation, a visiting professor of Computer Science at the Open University and an MIT Media Lab Research Affiliate.*

All opinions expressed by commentators are solely their own and do not reflect the opinions of Locus.