

AI is a bubble and it will burst. Most of the companies will fail. Most of the data-centers will be shuttered or sold for parts.

"We need to protect artists from AI predation, not just create a new way for artists to be mad about their impoverishment."



What will be left behind?

The **REVERSE**
CENTAUR'S
guide to
CRITICIZING AI

CORY DOCTOROW

All that stuff is terrible and harmful, but it's not driving investment. The total dollar figure represented by these apps doesn't come close to making a dent in the capital expenditures and operating costs of AI. They are peripheral, residual uses: flashy, but unimportant to the bubble.

Get rid of all those uses and you reduce the expected income of AI companies by a sum so small it rounds to zero.

Same goes for all that "AI Safety" nonsense, that purports to concern itself with preventing an AI from attaining sentience and turning us all into paperclips. First of all, this is facially absurd. Throwing more words and GPUs into the word-guessing program won't make it sentient. That's like saying, "Well, we keep breeding these horses to run faster and faster, so it's only a matter of time until one of our mares gives birth to a locomotive." A human mind is not a word-guessing program with a lot of extra words.

I'm here for science fiction thought experiments, don't get me wrong. But also, don't mistake SF for prophecy. SF stories about superintelligence are futuristic parables, not business plans, roadmaps, or predictions.

The AI Safety people say they are worried that AI is going to end the world, but AI bosses *love* these weirdos. Because on the one hand, if AI is powerful enough to destroy the world, think of how much money it can make! And on the other hand, no AI business plan has a line on its revenue projections spreadsheet labeled "Income from turning the human race into paperclips." So even if we ban AI companies from doing this, we won't cost them a dime in investment capital.

To pop the bubble, we have to hammer on the forces that created the bubble: the myth that AI can do your job, especially if you get high wages that your boss can claw back; the understanding that growth companies *need* a succession of ever-more-outlandish bubbles to stay alive; the fact that workers and the public they serve are on one side of this fight, and bosses and their investors are on the other side.

Because the AI bubble really *is* very bad news, it's worth fighting seriously, and a serious fight against AI strikes at its roots: the material factors fueling the hundreds of billions in wasted capital that are being spent to put us all on the breadline and fill all our walls with high-tech asbestos.

This text is from a talk given by Cory Doctorow at the University of Washington on December 4, 2025, as part of the "Neuroscience, AI and Society" lecture series, through the university's Computational Neuroscience Center.

Cory Doctorow (craphound.com) is a science fiction author, activist and journalist. He is the author of dozens of books, most recently *Enshittification: Why Everything Suddenly Got Worse And What To Do About It*. In 2020, he was inducted into the Canadian Science Fiction and Fantasy Hall of Fame.

This work is openly licensed via CC BY 4.0
<https://creativecommons.org/licenses/by/4.0/>

Cover image: *Pursued Centaur* (1881) by Max Klinger
from The Art Institute of Chicago's Public Domain Collection

AI is a bubble and it will burst. Most of the companies will fail. Most of the data-centers will be shuttered or sold for parts. So what will be left behind?

We'll have a bunch of coders who are really good at applied statistics. We'll have a lot of cheap GPUs, which'll be good news for, say, effects artists and climate scientists, who'll be able to buy that critical hardware at pennies on the dollar. And we'll have the open source models that run on commodity hardware, AI tools that can do a lot of useful stuff, like transcribing audio and video, describing images, summarizing documents, automating a lot of labor-intensive graphic editing, like removing backgrounds, or airbrushing passersby out of photos. These will run on our laptops and phones, and open source hackers will find ways to push them to do things their makers never dreamt of.

If there had never been an AI bubble, if all this stuff arose merely because computer scientists and product managers noodled around for a few years coming up with cool new apps for back-propagation, machine learning and generative adversarial networks, most people would have been pleasantly surprised with these interesting new things their computers could do. We'd call them "plugins."

It's the bubble that sucks, not these applications. The bubble doesn't want cheap useful things. It wants expensive, "disruptive" things: Big foundation models that lose billions of dollars every year.

When the AI investment mania halts, most of those models are going to disappear, because it just won't be economical to keep the data-centers running. As Stein's Law has it: "Anything that can't go on forever eventually stops."

The collapse of the AI bubble is going to be ugly. Seven AI companies currently account for more than a third of the stock market, and they endlessly pass around the same \$100B IOU.

Bosses are mass-firing productive workers and replacing them with janky AI, and when the janky AI is gone, no one will be able to find and re-hire most of those workers, we're going to go from dysfunctional AI systems to nothing.

AI is the asbestos in the walls of our technological society, stuffed there with wild abandon by a finance sector and tech monopolists run amok. We will be excavating it for a generation or more.

So we need to get rid of this bubble. Pop it, as quickly as we can. To do that, we have to focus on the material factors driving the bubble. The bubble isn't being driven by deepfake porn, or election disinformation, or AI image-gen, or slop advertising.

I'm a science fiction writer, which means that my job is to make up futuristic parables about our current techno-social arrangements to interrogate not just what a gadget *does*, but who it does it *for*, and who it does it *to*.

What I *don't* do is predict the future. *No one* can predict the future, which is a good thing, since if the future were predictable, that would mean that what we all do couldn't change it. It would mean that the future was arriving on fixed rails and couldn't be steered.

Jesus Christ, what a miserable proposition!

Now, not everyone understands the distinction. They think SF writers are oracles, soothsayers. Unfortunately, even some of my colleagues labor under the delusion that they can "see the future."

But for every SF writer who deludes themselves into thinking that they are writing the future, there are a hundred SF fans who believe that they are *reading* the future, and a depressing number of those people appear to have become AI bros. The fact that these guys can't shut up about the day that their spicy autocomplete machine will wake up and turn us all into paperclips has led many confused journalists and conference organizers to try to get me to comment on the future of AI.

That's a thing I strenuously resisted doing, because I wasted two years of my life explaining patiently and repeatedly why I thought crypto was stupid, and getting relentless bollocked by cryptocurrency cultists who at first insisted that I just didn't understand crypto. And then, when I made it clear that I *did* understand crypto, insisted that I must be a paid shill.

This is literally what happens when you argue with Scientologists, and life is Just. Too. Short.

So I didn't want to get lured into another one of those quagmires, because on the one hand, I just don't think AI is that important of a technology, and on the other hand, I have very nuanced and complicated views about what's wrong, and not wrong, about AI, and it takes a long time to explain that stuff.

But people wouldn't stop asking, so I did what I always do. I wrote a book.

Over the summer I wrote a book about what I think about AI, which is really about what I think about AI criticism, and more specifically, how to be a good AI critic. By which I mean: "How to be a critic whose criticism inflicts maximum damage on the parts of AI that are doing the most harm. I titled the book *The Reverse Centaur's Guide to Life After AI*, and Farrar, Straus and Giroux will publish it in June, 2026.

But you don't have to wait until then because I am going to break down the entire book's thesis for you tonight, over the next 40 minutes. I am going to talk *fast*.



Start with what a reverse centaur is. In automation theory, a "centaur" is a person who is assisted by a machine. You're a human head being carried around on a tireless robot body. Driving a car makes you a centaur, and so does using autocomplete.

And obviously, a *reverse* centaur is machine head on a human body, a person who is serving as a squishy meat appendage for an uncaring machine.

Like an Amazon delivery driver, who sits in a cabin surrounded by AI cameras, that monitor the driver's eyes and take points off if the driver looks in a proscribed direction, and monitors the driver's mouth because singing isn't allowed on the job, and rats the driver out to the boss if they don't make quota.

The driver is in that van because the van can't drive itself and can't get a parcel from the curb to your porch. The driver is a peripheral for a van, and the van drives the driver, at superhuman speed, demanding superhuman endurance. But the driver is human, so the van doesn't just use the driver. The van uses the driver *up*.

Obviously, it's nice to be a centaur, and it's horrible to be a reverse centaur. There are lots of AI tools that are potentially very centaur-like, but my thesis is that these tools are created and funded for the express purpose of creating reverse-centaurs, which is something none of us want to be.

But like I said, the job of an SF writer is to do more than think about what the gadget does, and drill down on who the gadget does it *for* and who the gadget does it *to*. Tech bosses want us to believe that there is only one way a technology can be used. Mark Zuckerberg wants you to think that it's technologically impossible to have a conversation with a friend without him listening in. Tim Cook wants you to think that it's technologically impossible for you to have a reliable computing experience unless he gets a veto over which software you install and without him taking 30 cents out of every dollar you spend. Sundar Pichai wants you think that it's impossible for you to find a webpage unless he gets to spy on you from asshole to appetite.

This is all a kind of vulgar Thatcherism. Margaret Thatcher's mantra was "There is no alternative." She repeated this so often they called her "TINA" Thatcher: There. Is. No. Alternative. TINA.

with ideas or variations, that's no problem, because the ultimate work comes from you. And if you're a video editor who uses deepfakes to change the eyelines of 200 extras in a crowd-scene, then sure, those eyeballs are in the public domain, but the movie stays copyrighted.

But creative workers don't have to rely on the US government to rescue us from AI predators. We can do it ourselves, the way the writers did in their historic writers' strike. The writers brought the studios to their knees. They did it because they are organized and solidaristic, but also are allowed to do something that virtually no other workers are allowed to do: they can engage in "sectoral bargaining," whereby all the workers in a sector can negotiate a contract with every employer in the sector.

That's been illegal for most workers since the late 1940s, when the Taft-Hartley Act outlawed it. If we are gonna campaign to get a new law passed in hopes of making more money and having more control over our labor, we should campaign to restore sectoral bargaining, not to expand copyright.

Our allies in a campaign to expand copyright are our bosses, who have never had our best interests at heart. While our allies in the fight for sector bargaining are every worker in the country. As the song goes, "Which side are you on?"

OK, I need to bring this talk in for a landing now, because I'm out of time, so I'm going to close out with this: AI is a bubble and bubbles are terrible.

Bubbles transfer the life's savings of normal people who are just trying to have a dignified retirement to the wealthiest and most unethical people in our society, and every bubble eventually bursts, taking their savings with it.

But not every bubble is created equal. Some bubbles leave behind something productive. Worldcom stole billions from everyday people by defrauding them about orders for fiber optic cables. The CEO went to prison and died there. But the fiber outlived him. It's still in the ground. At my home, I've got 2gb symmetrical fiber, because AT&T lit up some of that old Worldcom dark fiber.

All things being equal, it would have been better if Worldcom hadn't ever existed, but the only thing worse than Worldcom committing all that ghastly fraud would be if there was nothing to salvage from the wreckage.

I don't think we'll salvage much from cryptocurrency, for example. Sure, there'll be a few coders who've learned something about secure programming in Rust. But when crypto dies, what it will leave behind is bad Austrian economics and worse monkey JPEGs.

These companies are *desperate* to use AI to displace workers. When Getty Images sues AI companies, it's not representing the interests of *photographers*. Getty *hates* paying photographers! Getty just wants to get paid for the training run, and they want the resulting AI model to have guardrails, so it will refuse to create images that compete with Getty's images for anyone except Getty. But Getty will *absolutely* use its models to bankrupt as many photographers as it possibly can.

A new copyright to train models won't get us a world where models aren't used to destroy artists, it'll just get us a world where the standard contracts of the handful of companies that control all creative labor markets are updated to require us to hand over those new training rights to those companies. Demanding a new copyright just makes you a useful idiot for your boss, a human shield they can brandish in policy fights, a tissue-thin pretense of "won't someone think of the hungry artists?"

When really what they're demanding is a world where 30% of the investment capital of the AI companies go into their shareholders' pockets. When an artist is being devoured by rapacious monopolies, does it matter how they divvy up the meal?

We need to protect artists from AI predation, not just create a new way for artists to be mad about their impoverishment.

And incredibly enough, there's a really simple way to do that. After 20+ years of being consistently wrong and terrible for artists' rights, the US Copyright Office has finally done something gloriously, wonderfully *right*. All through this AI bubble, the Copyright Office has maintained – correctly – that AI-generated works cannot be copyrighted, because copyright is exclusively for humans. That's why the "monkey selfie" is in the public domain. Copyright is only awarded to works of human creative expression that are fixed in a tangible medium.

And not only has the Copyright Office taken this position, they've defended it vigorously in court, repeatedly winning judgments to uphold this principle.

The fact that every AI created work is in the public domain means that if Getty or Disney or Universal or Hearst newspapers use AI to generate works – then anyone else can take those works, copy them, sell them, or give them away for free. And the only thing those companies hate more than paying creative workers, is having other people take *their* stuff without permission.

The US Copyright Office's position means that the only way these companies can get a copyright is to *pay humans* to do creative work. This is a recipe for centaurhood. If you're a visual artist or writer who uses prompts to come up

"There is no alternative" is a cheap rhetorical slight. It's a demand dressed up as an observation. "There is no alternative" means "STOP TRYING TO THINK OF AN ALTERNATIVE." Which, you know, *fuck that*.

I'm a SF writer, my job is to think of a dozen alternatives before breakfast.

So let me explain what I think is going on here with this AI bubble, and sort out the bullshit from the material reality, and explain how I think we could and should all be better AI critics.



Start with monopolies: tech companies are gigantic and they don't compete, they just take over whole sectors, either on their own or in cartels.

Google and Meta control the ad market. Google and Apple control the mobile market, and Google pays Apple more than \$20 billion/year not to make a competing search engine, and of course, Google has a 90% Search market-share.

Now, you'd think that this was good news for the tech companies, owning their whole sector.

But it's actually a crisis. You see, when a company is growing, it is a "growth stock," and investors really like growth stocks. When you buy a share in a growth stock, you're making a bet that it will continue to grow. So growth stocks trade at a huge multiple of their earnings. This is called the "price to earnings ratio" or "P/E ratio."

But once a company stops growing, it is a "mature" stock, and it trades at a much *lower* P/E ratio. So for every dollar that Target – a mature company – brings in, it is worth ten dollars. It has a P/E ratio of 10, while Amazon has a P/E ratio of 36, which means that for every dollar Amazon brings in, the market values it at \$36.

It's wonderful to run a company that's got a growth stock. Your shares are as good as money. If you want to buy another company, or hire a key worker, you can offer stock instead of cash. And stock is very easy for companies to get, because shares are manufactured right there on the premises, all you have to do is type some zeroes into a spreadsheet, while dollars are much harder to come by. A company can only get dollars from customers or creditors.

So when Amazon bids against Target for a key acquisition, or a key hire, Amazon can bid with shares they make by typing zeroes into a spreadsheet, and Target can only bid with dollars they get from selling stuff to us, or taking out loans, which is why Amazon generally wins those bidding wars.

That's the *upside* of having a growth stock. But here's the downside: eventually a company has to stop growing. Like, say you get a 90% market share in your sector, how are you gonna grow?

Once the market decides that you aren't a growth stock, once you become mature, your stock is revalued, to a P/E ratio befitting a mature stock.

If you are an exec at a dominant company with a growth stock, you have to live in constant fear that the market will decide that you're not likely to grow any further. Think of what happened to Facebook in the first quarter of 2022. They told investors that they experienced slightly slower growth in the USA than they had anticipated, and investors panicked. They staged a one-day, \$240B sell off. A quarter-trillion dollars in 24 hours! At the time, it was the largest, most precipitous drop in corporate valuation in human history.

That's a monopolist's worst nightmare, because once you're presiding over a "mature" firm, the key employees you've been compensating with stock, experience a precipitous pay-drop and bolt for the exits, so you lose the people who might help you grow again, and you can only hire their replacements with dollars. With dollars, not shares.

And the same goes for acquiring companies that might help you grow, because they, too, are going to expect money, not stock. This is the paradox of the growth stock. While you are growing to domination, the market loves you, but once you *achieve* dominance, the market lops 75% or more off your value in a single stroke if they don't trust your pricing power.

Which is why growth stock companies are always desperately pumping up one bubble or another, spending billions to hype the pivot to video, or cryptocurrency, or NFTs, or Metaverse, or AI.

I'm not saying that tech bosses are making bets they don't plan on winning. But I am saying that winning the bet – creating a viable metaverse – is the secondary goal. The primary goal is to keep the market convinced that your company will continue to grow, and to remain convinced until the next bubble comes along.

So this is *why* they're hyping AI: the material basis for the hundreds of billions in AI investment.



Now I want to talk about *how* they're selling AI. The growth narrative of AI is that AI will disrupt labor markets. I use "disrupt" here in its most disreputable, tech bro sense.

It ends: "This action by Disney and Universal represents a critical stand for human creativity and responsible innovation."

And it's signed by Mitch Glazier, CEO of the RIAA.

It's very likely that name doesn't mean anything to you. But let me tell you who Mitch Glazier is. Today, Mitch Glazier is the CEO of the RIAA, with an annual salary of \$1.3M. But until 1999, Mitch Glazier was a key Congressional staffer, and in 1999, Glazier snuck an amendment into an unrelated bill, the Satellite Home Viewer Improvement Act, that killed musicians' right to take their recordings back from their labels.

This is a practice that had been especially important to "heritage acts" (which is a record industry euphemism for "old music recorded by Black people"), for whom this right represented the difference between making rent and ending up on the street.

When it became clear that Glazier had pulled this musician-impoverishing scam, there was so much public outcry, that Congress actually came back for a special session, just to vote again to cancel Glazier's amendment. And then Glazier was kicked out of his cushy Congressional job, whereupon the RIAA started paying him more than \$1M/year to "represent the music industry."

This is the guy who signed that press release in my inbox. And his message was: *The problem isn't that Midjourney wants to train a Gen AI model on copyrighted works, and then use that model to put artists on the breadline. The problem is that Midjourney didn't pay RIAA members Universal and Disney for permission to train a model. Because if only Midjourney had given Disney and Universal several million dollars for training rights to their catalogs, the companies would have happily allowed them to train to their heart's content, and they would have bought the resulting models, and fired as many creative professionals as they could.*

I mean, have we already forgotten the Hollywood strikes? I sure haven't. I live in Burbank, home to Disney, Universal and Warner, and I was out on the line with my comrades from the Writers Guild, offering solidarity on behalf of my union, IATSE 830, The Animation Guild, where I'm a member of the writers' unit.

And I'll never forget when one writer turned to me and said, "You know, you prompt an LLM exactly the same way an exec gives shitty notes to a writers' room. You know: 'Make me ET, except it's about a dog, and put a love interest in there, and a car chase in the second act.' The difference is, you say that to a writers' room and they all make fun of you and call you a fucking idiot suit. But you say it to an LLM and it will cheerfully shit out a terrible script that conforms exactly to that spec (you know, *Air Bud*)."

But hey, maybe you think this is all sophistry. Maybe you think I'm full of shit. That's fine. It wouldn't be the first time someone thought that.

After all, even if I'm right about how copyright works now, there's no reason we couldn't change copyright to ban training activities, and maybe there's even a clever way to wordsmith the law so that it only catches bad things we don't like, and not all the good stuff that comes from scraping, analyzing and publishing.

Well, even then, you're not gonna help out creators by creating this new copyright. If you're thinking that you can, you need to grapple with this fact: we have monotonically expanded copyright since 1976, so that today, copyright covers more kinds of works, grants exclusive rights over more uses, and lasts longer.

And today, the media industry is larger and more profitable than it has ever been, and also: the share of media industry income that goes to creative workers is lower than its ever been, both in real terms, and as a proportion of those incredible gains made by creators' bosses at the media company.

So how it is that we have given all these new rights to creators, and those new rights have generated untold billions, and left creators *poorer*? It's because in a creative market dominated by five publishers, four studios, three labels, two mobile app stores, and a single company that controls all the ebooks and audiobooks, giving a creative worker extra rights to bargain with is like giving your bullied kid more lunch money.

It doesn't matter how much lunch money you give the kid, the bullies will take it all. Give that kid enough money and the bullies will hire an agency to run a global campaign proclaiming "think of the hungry kids! Give them more lunch money!"

Creative workers who cheer on lawsuits by the big studios and labels need to remember the first rule of class warfare: things that are good for your boss are rarely what's good for you.

The day Disney and Universal filed suit against Midjourney, I got a press release from the RIAA, which represents Disney and Universal through their recording arms. Universal is the largest label in the world. Together with Sony and Warner, they control 70% of all music recordings in copyright today.

It starts: "There is a clear path forward through partnerships that both further AI innovation and foster human artistry."

The promise of AI – the promise AI companies make to investors – is that there will be AIs that can do your job, and when your boss fires you and replaces you with AI, he will keep half of your salary for himself, and give the other half to the AI company.

That's it.

That's the \$13T growth story that MorganStanley is telling. It's why big investors and institutionals are giving AI companies hundreds of billions of dollars. And because they are piling in, normies are also getting sucked in, risking their retirement savings and their family's financial security.

Now, if AI could do your job, this would *still* be a problem. We'd have to figure out what to do with all these technologically unemployed people.

But AI can't do your job. It can *help* you do your job, but that doesn't mean it's going to save anyone money. Take radiology: there's some evidence that AIs can sometimes identify solid-mass tumors that some radiologists miss, and look, I've got cancer. Thankfully, it's very treatable, but I've got an interest in radiology being as reliable and accurate as possible.

If my Kaiser hospital bought some AI radiology tools and told its radiologists: "Hey folks, here's the deal. Today, you're processing about 100 x-rays per day. From now on, we're going to get an instantaneous second opinion from the AI, and if the AI thinks you've missed a tumor, we want you to go back and have another look, even if that means you're only processing 98 x-rays per day. That's fine, we just care about finding all those tumors."

If that's what they said, I'd be delighted. But no one is investing hundreds of billions in AI companies because they think AI will make radiology more expensive, not even if that also makes radiology more accurate. The market's bet on AI is that an AI salesman will visit the CEO of Kaiser and make this pitch: "Look, you fire 9/10s of your radiologists, saving \$20M/year, you give us \$10M/year, and you net \$10M/year, and the remaining radiologists' job will be to oversee the diagnoses the AI makes at superhuman speed, and somehow remain vigilant as they do so, despite the fact that the AI is usually right, except when it's catastrophically wrong.

"And if the AI misses a tumor, this will be the human *radiologist's* fault, because they are the 'human in the loop.' It's their signature on the diagnosis."

This is a reverse centaur, and it's a specific kind of reverse-centaur: it's what Dan Davies calls an "accountability sink." The radiologist's job isn't really to oversee the AI's work, it's to take the blame for the AI's mistakes.

This is another key to understanding – and thus deflating – the AI bubble. The AI can't do your job, but an AI salesman can convince your boss to fire you and replace you with an AI that *can't* do your job. This is key because it helps us build the kinds of coalitions that will be successful in the fight against the AI bubble.

If you're someone who's worried about cancer, and you're being told that the price of making radiology too cheap to meter, is that we're going to have to re-home America's 32,000 radiologists, with the trade-off that no one will ever be denied radiology services again, you might say, "Well, OK, I'm sorry for those radiologists, and I fully support getting them job training or UBI or whatever. But the point of radiology is to fight cancer, not to pay radiologists, so I know what side I'm on."

AI hucksters and their customers in the C-suites want the public on their side. They want to forge a class alliance between AI deployers and the people who enjoy the fruits of the reverse centaurs' labor. They want us to think of ourselves as enemies to the workers.

Now, some people will be on the workers' side because of politics or aesthetics. They just like workers better than their bosses. But if you want to win over *all* the people who benefit from your labor, you need to understand and stress how the products of the AI will be substandard. That they are going to get charged more for worse things. That they have a shared material interest with you.

Will those products be substandard? There's every reason to think so. Earlier, I alluded to "automation blindness, "the physical impossibility of remaining vigilant for things that rarely occur. This is why TSA agents are incredibly good at spotting water bottles. Because they get a ton of practice at this, all day, every day. And why they fail to spot the guns and bombs that government red teams smuggle through checkpoints to see how well they work, because they just don't have any practice at that. Because, to a first approximation, no one deliberately brings a gun or a bomb through a TSA checkpoint.

Automation blindness is the Achilles' heel of "humans in the loop."

Think of AI software generation: there are plenty of coders who love using AI, and almost without exception, they are senior, experienced coders, who get to decide how they will use these tools. For example, you might ask the AI to generate a set of CSS files to faithfully render a web-page across multiple versions of multiple browsers. This is a notoriously fiddly thing to do, and it's pretty easy to verify if the code works – just eyeball it in a bunch of browsers. Or maybe the coder has a single data file they need to import and they don't want to write a whole utility to convert it.

And I'm here to tell you *they are wrong*: wrong because this would inflict terrible collateral damage on socially beneficial activities, and it would represent a massive expansion of copyright over activities that are currently permitted – for good reason!

Let's break down the steps in AI training.

First, you scrape a bunch of web-pages. This is unambiguously legal under present copyright law. You do *not* need a license to make a transient copy of a copyrighted work in order to analyze it, otherwise search engines would be illegal. Ban scraping and Google will be the last search engine we ever get, the Internet Archive will go out of business, that guy in Austria who scraped all the grocery store sites and proved that the big chains were colluding to rig prices would be in deep trouble.

Next, you perform analysis on those works. Basically, you count stuff on them: count pixels and their colors and proximity to other pixels; or count words. This is obviously not something you need a license for. It's just not illegal to count the elements of a copyrighted work. And we really don't want it to be, not if you're interested in scholarship of any kind.

And it's important to note that counting things is legal, even if you're working from an illegally obtained copy. Like, if you go to the flea market, and you buy a bootleg music CD, and you take it home and you make a list of all the adverbs in the lyrics, and you publish that list, you are not infringing copyright by doing so.

Perhaps you've infringed copyright by getting the pirated CD, but not by counting the lyrics.

This is why Anthropic offered a \$1.5B settlement for training its models based on a ton of books it downloaded from a pirate site: not because counting the words in the books infringes anyone's rights, but because they were worried that they were going to get hit with \$150K/book statutory damages for *downloading* the files.

OK, after you count all the pixels or the words, it's time for the final step: publishing them. Because that's what a model is: a literary work (that is, a piece of software) that embodies a bunch of facts about a bunch of other works, word and pixel distribution information, encoded in a multidimensional array.

And again, copyright absolutely does not prohibit you from publishing facts about copyrighted works. And again, no one should want to live in a world where someone else gets to decide which truthful, factual statements you can publish.

that is good art – as opposed to merely striking, or interesting, or an example of good draftsmanship – it will be thanks to those additional infusions of creative intent by a human.

And in the meantime, it's bad art. It's bad art in the sense of being "eerie," the word Mark Fisher uses to describe "when there is something present where there should be nothing, or there is nothing present when there should be something."

AI art is eerie because it seems like there is an intender and an intention behind every word and every pixel, because we have a lifetime of experience that tells us that paintings have painters, and writing has writers. But it's missing something. It has nothing to say, or whatever it has to say is so diluted that it's undetectable.

The images were striking before we figured out the trick, but now they're just like the images we imagine in clouds or piles of leaves. *We're* the ones drawing a frame around part of the scene, we're the ones focusing on some contours and ignoring the others. We're looking at an inkblot, and it's not telling us anything.

Sometimes that can be visually arresting, and to the extent that it amuses people in a community of prompters and viewers, that's harmless.

I know someone who plays a weekly Dungeons and Dragons game over Zoom. It's transcribed by an open source model running locally on the dungeon master's computer, which summarizes the night's session and prompts an image generator to create illustrations of key moments. These summaries and images are hilarious *because* they're full of errors. It's a bit of harmless fun, and it brings a small amount of additional pleasure to a small group of people. No one is going to fire an illustrator because D&D players are image-generating funny illustrations where seven-fingered paladins wrestle with orcs that have an extra hand.

But bosses have and *will* fire illustrators, because they fantasize about being able to dispense with creative professionals and just prompt an AI. Because even though the AI can't do the illustrator's job, an AI salesman can convince the illustrator's boss to fire them and replace them with an AI that can't do their job.

This is a disgusting and terrible juncture, and we should not simply shrug our shoulders and accept Thatcherism's fatalism: "There is no alternative."

So what *is* the alternative? A lot of artists and their allies think they have an answer: they say we should extend copyright to cover the activities associated with training a model.

Tasks like these can genuinely make coders more efficient and give them more time to do the fun part of coding, namely, solving really gnarly, abstract puzzles. But when you listen to business leaders talk about their AI plans for coders, it's clear they're not looking to make some centaurs.

They want to fire a *lot* of tech workers – 500,000 over the past three years – and make the rest pick up their work with coding, which is only possible if you let the AI do all the gnarly, creative problem solving, and then you do the most boring, soul-crushing part of the job: reviewing the AIs' code.

And because AI is just a word guessing program, because all it does is calculate the most probable word to go next, the errors it makes are especially subtle and hard to spot, because these bugs are literally statistically indistinguishable from working code (except that they're bugs).

Here's an example: code libraries are standard utilities that programmers can incorporate into their apps, so they don't have to do a bunch of repetitive programming. Like, if you want to process some text, you'll use a standard library. If it's an HTML file, that library might be called something like `lib.html.text.parsing`; and if it's a DOCX file, it'll be `lib.docx.text.parsing`. But reality is messy, humans are inattentive and stuff goes wrong, so sometimes, there's another library, this one for parsing PDFs, and instead of being called `lib.pdf.text.parsing`, it's called `lib.text.pdf.parsing`.

Now, because the AI is a statistical inference engine, because all it can do is predict what word will come next based on all the words that have been typed in the past, it will "hallucinate" a library called `lib.pdf.text.parsing`. And the thing is, malicious hackers *know* that the AI will make this error, so they will go out and *create* a library with the predictable, hallucinated name, and that library will get automatically sucked into your program, and it will do things like steal user data or try and penetrate other computers on the same network.

And you, the human in the loop – the reverse centaur – you have to spot this subtle, hard to find error, this bug that is literally statistically indistinguishable from correct code. Now, maybe a senior coder could catch this, because they've been around the block a few times, and they know about this tripwire.

But guess who tech bosses want to preferentially fire and replace with AI? Senior coders. Those mouthy, entitled, *extremely highly paid* workers, who don't think of themselves as workers. Who see themselves as founders in waiting, peers of the company's top management. The kind of coder who'd lead a walkout over the company building drone-targeting systems for the Pentagon, which cost Google *ten billion dollars* in 2018.

For AI to be valuable, it has to replace *high-wage* workers, and those are precisely the experienced workers, with process knowledge, and hard-won intuition, who might spot some of those statistically camouflaged AI errors.

Like I said, the point here is to replace high-waged workers.

And one of the reasons the AI companies are so anxious to fire coders is that coders are the princes of labor. They're the most consistently privileged, sought-after, and well-compensated workers in the labor force.

If you can replace *coders* with AI, who *can't* you replace with AI? Firing coders is an *ad* for AI.

Which brings me to AI *art*. AI art – or “art” – is also an ad for AI, but it's not part of AI's business model.

Let me explain: on average, illustrators don't make any money. They are already one of the most immiserated, precartized groups of workers out there. They suffer from a pathology called “vocational awe.” That's a term coined by the librarian Fobazi Ettarh, and it refers to workers who are vulnerable to workplace exploitation because they actually care about their jobs – nurses, librarians, teachers, and artists.

If AI image generators put every illustrator working today out of a job, the resulting wage-bill savings would be undetectable as a proportion of all the costs associated with training and operating image-generators. The total wage bill for commercial illustrators is less than the kombucha bill for the company cafeteria at just *one* of Open AI's campuses.

The purpose of AI art – and the story of AI art as a death-knell for artists – is to convince the broad public that AI is amazing and will do amazing things. It's to create buzz. Which is not to say that it's not disgusting that former OpenAI CTO Mira Murati told a conference audience that “some creative jobs shouldn't have been there in the first place,” and that it's not especially disgusting that she and her colleagues boast about using the work of artists to ruin those artists' livelihoods.

It's *supposed* to be disgusting. It's supposed to get artists to run around and say, “The AI can do my job, and it's going to steal my job, and isn't that terrible?”

Because the customers for AI – corporate bosses – don't see AI taking workers' jobs as terrible. They see it as wonderful.

But can AI do an illustrator's job? Or any artist's job?

Let's think about that for a second. I've been a working artist since I was 17 years old, when I sold my first short story, and I've given it a lot of thought, and here's what I think art is: it starts with an artist, who has some vast, complex, numinous, irreducible feeling in their mind. And the artist infuses that feeling into some artistic medium. They make a song, or a poem, or a painting, or a drawing, or a dance, or a book, or a photograph. And the idea is, when *you* experience this work, a facsimile of the big, numinous, irreducible feeling will materialize in *your* mind.

Now that I've defined art, we have to go on a little detour.

I have a friend who's a law professor, and before the rise of chatbots, law students knew better than to ask for reference letters from their profs, unless they were a *really* good student. Because those letters were a pain in the ass to write. So if you advertised for a postdoc and you heard from a candidate with a reference letter from a respected prof, the mere existence of that letter told you that the prof really thought highly of that student.

But then we got chatbots, and everyone knows that you generate a reference letter by feeding three bullet points to an LLM, and it'll barf up five paragraphs of florid nonsense about the student.

So when my friend advertises for a postdoc, they are *flooded* with reference letters, and they deal with this flood by feeding all these letters to *another* chatbot, and ask it to reduce them back to three bullet points. Now, obviously, they won't be the same bullet-points, which makes this whole thing terrible.

But just as obviously, nothing in that five-paragraph letter *except* the original three bullet points are relevant to the student. The chatbot doesn't know the student. It doesn't know anything about them. It cannot add a single true or useful statement about the student to the letter.

What does this have to do with AI art? Art is a transfer of a big, numinous, irreducible feeling from an artist to someone else. But the image-gen program doesn't know *anything* about your big, numinous, irreducible feeling. The only thing it knows is whatever you put into your prompt, and those few sentences are diluted across a million pixels or a hundred thousand words, so that the average communicative density of the resulting work is indistinguishable from zero.

It's possible to infuse more communicative intent into a work: writing more detailed prompts, or doing the selective work of choosing from among many variants, or directly tinkering with the AI image after the fact, with a paintbrush or Photoshop or The Gimp. And if there will ever be a piece of AI art