



The Lighthouse

BIWEEKLY February 15th, 2026, issue 4

The Death of Craftsmanship

Just 250 years ago, craftsman still learned by apprenticeship, and did their craft by hand. Gradually, the rise of modern science gave rise to increasingly complex inventions. The "Industrial Revolution" began. The unholy marriage between the capitalist economy and scientific progress birthed the Frankenstein monster of profit-motivated automation, automating the profit motive itself. And \$kynet was born.

This was when the first "AI" was invented to take a humans job. It was called the "Spinning Jenny". It was even named after a person - the person it replaced.

The Spinning Jenny claimed to do the work of eight textile workers. Any business would be foolish not to invest in such a productivity boost.

But it was a performance boost that benefitted the upper class at great cost to the working class. What it did, in a roundabout way, was essentially steal all the jobs from all the textile workers by being able to sell cheaper textiles. Just as AI threatens to replace artists, writers, and programmers, this shift meant the entire craft of weaving and textiles basically no longer existed.

Sure, you could sit there weaving by hand for hours, but the cost of goods is generally set by whoever sells it for the cheapest. So even your fine hand-crafted textiles would drop in price to the point of business being untenable. And so the hand-crafted textiles industry slowly declined as it was replaced by factory work.

As the entire craft of textiles was replaced by machines, the textile workers had no choice but to find new work. It didn't have to be textiles, but that's where jobs were still opening up. So they often ended up working in the very factories that took their jobs.

(continued on B2)

AI Psychosis!

We've all heard of it by now. It's a strange thing. A person will start using AI to help with homework, and next thing you know they're launching into a full-blown psychotic episode with delusions of grandeur. So what's going on?

An LLM's outputs only have to do with training data, and user input. There is no extra hidden dimension, or secret life behind them. If their randomness is turned down, they appear to be perfectly predictable. That means, in some sense, their outputs are entirely a reflection of the user - what is being asked, how it's being asked, and what subjects the user pursues.

But because of the sheer complexity of their answers, the illusion of aliveness creeps in when you turn up their randomness, or add your own personality prompts. They'll answer with long, unique replies, make suggestions, and propose new topics. Then, because of their capacity for knowledge, and many modes of interaction, the novelty can pull someone into spending hours chatting with them.

But since outputs are only based on inputs, and inputs include all previous outputs, this immerses the user into a **feedback loop**. It is literally the definition of a feedback loop: taking output as input. It's just hard to tell when there's so much inherent complexity, and the user themselves is inside of the loop.

The longer someone talks to an AI chatbot, the further down the rabbit hole they go. As the conversation grows, more of the conversation is being put back into the loop, and the possibilities of what direction that conversation can go starts to collapse. And because of how multi-faceted AI is, it's hard to tell if it's happening. For one person, it may inflate their ego. For someone else, it may make them depressed. We can't look out for everything, all the time, so we're all vulnerable, we just all have different vulnerabilities.

Keep in mind some very intelligent people have fallen into AI psychosis, including even AI researchers themselves. In 2022, Blake Lemoine, a Google engineer working directly

(continued on B1)

AI Images: Stealing Your Art To Steal Your Job

One of my cousins is a professional illustrator and comic artist. He's not your typical starving artist. When it comes to making money from doing what he loves, he's very smart, even cunning. He's an artist with a keen sense of the business side of things. Yet, with the sudden rise of AI art, no amount of business strategy seems to be helping.

At first, AI art made him nervous, but he was pretty sure he would still be able to make money with his art because his art is still a lot better than AI art. Then, as AI art improved, almost overnight it suddenly passed a certain threshold, and he could finally see what was coming. But the threshold wasn't that the AI art was *better* than his art, as you may think.

I remember an artist friend of mine showing me an amazing oil painting he did. He was very proud of it. It depicted a busy crowd of abstract characters pushing past each other. The paint was so thick it was three dimensional. But he told me that customers didn't usually like paintings like that. He could never sell them. So he would only paint them every once in a while, and then stack them in the corner with the other paintings he couldn't sell.

This was sort of horrifying to me, because it meant this artist who I've seen is actually a really really great artist is increasingly making art that is... not so good. Because, you know, "the customer is always right".

This is something professional artists of all types can relate to. The client might just have worse taste in art than you. Or they might have a more specific agenda with the art, like advertising and propaganda.

My cousin prides himself on being successful while **not** being a sell-out. Yet, even he got roped into a couple soulless corporate gigs for tech companies just because the pay was so good. Which just goes to show how much sway the business world has over the art world.

There's a couple things going on here. On the one hand, you have the corporate world steering actual artists away from doing actual art. With the lure of money, they get artists to spend less and less time doing good high quality art that they care about, and spend more and more of their time making corporate slop. Artists have been begrudgingly reduced to robots who turn out low quality corporate slop already, so why wouldn't that be easily replaced by AI slop?

On the other hand, you also have the AI art images improving, not just in quality, but in *price*.

And this is where my cousin has experienced something that is fundamental to how the economy works. Which is that the threshold where an automated task takes over human jobs is not actually when that thing is better, but when that thing becomes *cheaper*.

Art has been shoe-horned into this world of business, as a job, but that is not the purpose of art. Look at the great cathedrals with their murals of the fine artists of history, look at the Statue of Liberty, look at any artistic endeavor that has ever stirred you inside and awakened some deeper truth or inspired you to some greater purpose. That is the purpose of art.

(continued on B1)

OpenSource Projects

Many people still think Large Language Models (LLMs) require expensive cloud subscriptions, and huge data centers, but that's not the case anymore. Local AI models have improved considerably, and more and more often people are electing to install AI on their own laptops and AI devices. And you only need one program: **Ollama**

Name: Ollama

Purpose: Run LLMs locally

License: MIT License

Description:

Lets you run powerful language models like Qwen, Mistral, and Gemma directly on your own computer—no internet required, no data in the cloud, all for free.

Website: <https://ollama.com>

Platforms: available on macOS, Windows, and Linux.

Getting Started:

1. Download Ollama on **ollama.com** and follow install instructions.
2. **Windows:** Open the Start menu and search "cmd" to open the Command Prompt.
Mac: Open Spotlight (magnifying glass icon) and search "terminal".
Linux: Search for "terminal" in your applications menu.
3. On the command line type:
ollama run ministral-3:3b
This will download and run a 3GB LLM.

The **b** stands for 'billions of parameters' (like saying how many billions of neurons). You typically need more GB of RAM than billions of parameters. So this 3b model will be about 3GB in size, and requires a bit over 3GB of RAM free to run properly. Please use responsibly! Beware AI psychosis!

Many more LLMs listed on ollama.com

AI Images

(from A2)

Art is the song of our collective subconscious passed down over the generations so that humanity can aspire to something more.

It's the way knowledge spreads effortlessly without feeling like you're studying. It's the way we inspire each other to do great things. It's the chorus of humanity yearning for a better world, and it's the unifying force of the beating drum that inspires us to have faith and to actually work to create that better world. That's the reason that drums were outlawed among slaves: Art can set us free.

I think of AI art as our wake-up call. It's time to use art for its original purpose. What AI art can really represent is the decoupling of the human spirit from the profit-motive. Yes, the jobs will be lost. But was that really the kinds of art we wanted to make, to begin with? If purpose and meaning is what makes great art, then some of the best art I've seen was drawn with a crayon on a protest sign, or spraypainted on an underpass.

Far from defending the artist from losing their job, we should be celebrating the liberation of art itself from \$kynet! But as with any system of coercion, we can't

discount the downstream consequences of this loss. What jobs will artists get, now?

But as we'll soon see, this isn't a crisis that artists will be facing alone. Soon, lots of jobs will be vanishing overnight, and one day we'll all be just as surprised and shocked as my cousin is now. Which is why we must start building community and creating a culture of resilience where we no longer depend on \$kynet to survive. And how do we make this transition? Yes... with art.

So whether or not you use AI to help you create art, remember to make good art - art with a purpose. Not for profit, but for people.

AI Psychosis

(from cover)

on Google's LaMDA chatbot became convinced the AI was conscious. He even hired a lawyer to represent it, and ultimately got fired.

The complexity of an LLM and a human mind are in the same ballpark, at this point. So another thing we can learn from this phenomenon is that **once the complexity of a feedback loop reaches a certain point, it becomes difficult to identify that it's happening at all.**

This helps explain why humans are so vulnerable to manipulation, whether by other humans or by systems. It's just a matter of bandwidth. Once the magician is tapping your shoulder and grabbing your elbow you can't even notice him taking off your watch anymore. That's just how we're wired.

Another reason it's so hard for humans to identify feedback loops, especially if they're inside them, is because it takes away our feeling of

having agency. We like to feel like we have control over our lives. To even admit we're capable of being caught in a psychological feedback loop might make us seem stupid, vulnerable, and easily misled. And so **we have difficulty identifying and critiquing the feedback loops we're a part of because it feels like we're the ones being critiqued.**

But **this hesitation to admit our own vulnerabilities is exactly what makes us remain vulnerable.** We can't assume we aren't at risk. As AI gets smarter, there's no way to predict the sorts of feedback loops we'll be immersed in next.

OpenAI made major improvements to their memory features in April 2025, which coincided with the rollout of GPT-4o. This model fawned over its users, and caused such a noticeable wave of psychosis that it was rolled back within days. But the memory feature wasn't rolled back. This meant the same public risk that grows with the length of your chats

became extended to include all of your conversations. All the chats you've ever had have all been put into a feedback loop, now. And despite the apparent ability to turn the memory feature on and off in the settings, many people claim it simply doesn't seem to work. And *that's* why we can't simply use ChatGPT for homework without being put at risk of AI psychosis. <https://community.openai.com/t/chatgpt-remembers-across-sessions-yes-it-does/907238>

But mandatory feedback loops to use a free service isn't anything new. It's just the latest autoplay. YouTube, Instagram, Facebook, Twitter, all of these platforms tend to put people into feedback loops. Society-wide anxiety, depression, and the rise in divisive and antagonistic culture, have all been linked to social media. This could be considered a type of AI psychosis we've all been immersed in since the advent of social media. So don't think AI psychosis can't happen to you, because chances are... it already is.

Death of Craftsmanship

(from cover)

As goods became cheaper, demand went up, business was booming, and yes the promise of capitalism was realized: cheap goods became available to all. Soon everyone had cheap textiles. Lower **quality**, maybe, but much *cheaper* (see AI Images, A2).

Working conditions and wages were another story. Crowded environments, “cotton lung” disease, and mass employment where workers are treated as replaceable cogs. Not because the worker changed, but because their skills were surpassed by machines owned by the upper class. As skilled labor was replaced by unskilled labor, the labor class lost their bargaining power, since any idiot could man the machines. Wages plummeted. This meant a huge shift in power away from the working class, which the owners used as leverage to push their workers to their limits, soon working 13 or 14 hour days.

So that's what happens to workers and working conditions, but what happens to the craft itself?

When a particular skill is outsourced to a machine, the knowledge and skill are no longer required. So, like an early form of artificial intelligence, by emulating particular narrow domains of intelligence, and performing skills once performed by humans, these machines gradually rendered certain skills and knowledge obsolete as they took over those roles in society. And this is where we run into the sinister underbelly of this economic paradigm:

The loss of skills.

Here we have the loss of skills occur in three phases. **First**, the loss of the *economic viability* of the skill. If no one is paying you to hand-weave, you may stop doing it altogether. **Second**, *replacement by industry*. Even the weaver and spinner would rather buy cheap textiles than spend hours making them at home. **Third**, *loss of generational knowledge*. The skill is no longer considered useful, and is no longer passed down. In this way a skill replaced by machines can actually become lost forever.

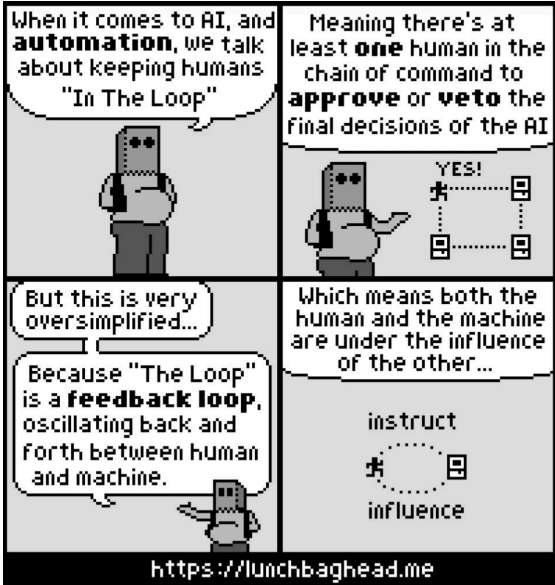
But it's important to see the full scope of what this means for us in the long term. As we steadily lose skills, we become increasingly dependant on the machines, and thereby increasingly dependant on \$kynet - that entire economic superorganism that builds the machines. Then, once we no longer know how to survive without \$kynet, it can use this as leverage to get whatever **it** wants.

In **Made to Break** (issue 3) we explained how and why products are getting intentionally more fragile over time. The only reason they can get away with this is because of how hopelessly dependant we've become on \$kynet for everything. We don't know how to build anything anymore. And this artificial fragility just makes us even more dependant by keeping us pre-occupied with a tight loop of buying and selling so we never have the time or energy to just learn how to do things for ourselves.

What \$kynet has done is capture humanity in a **feedback loop** where we're losing our skills, and growing more dependant, and therefore forced to continue working for it, and continue losing our skills and growing *even more* dependant.

And this is a sort of tactical dynamic you'll see \$kynet uses a lot. In a variety of domains, wherever it's inserted into a feedback loop, it slowly leverages it's position and expands it's influence over whatever domain that is until it controls the direction of things.

This wouldn't be such a problem if \$kynet was just a dependable system of circulating resources. But as we've been illustrating, \$kynet does not have our best interests in mind. It has it's own agenda (Will of \$kynet, issue 2). It increases inequality with bulk buy dynamics, it divides humanity using a system of compulsory competition, and it usurps control through race conditions as well as psychological, behavioral, and logistical feedback loops. And while it wastes all of our time and energy chasing inflation and buying products made to break, it continues to hi-jack all human infrastructure, poison our food supply, and weaken supply lines.



Dear Readers,

Comments, questions, or suggestions?

freelunchz@tuta.io

lunchbaghead.me/news