

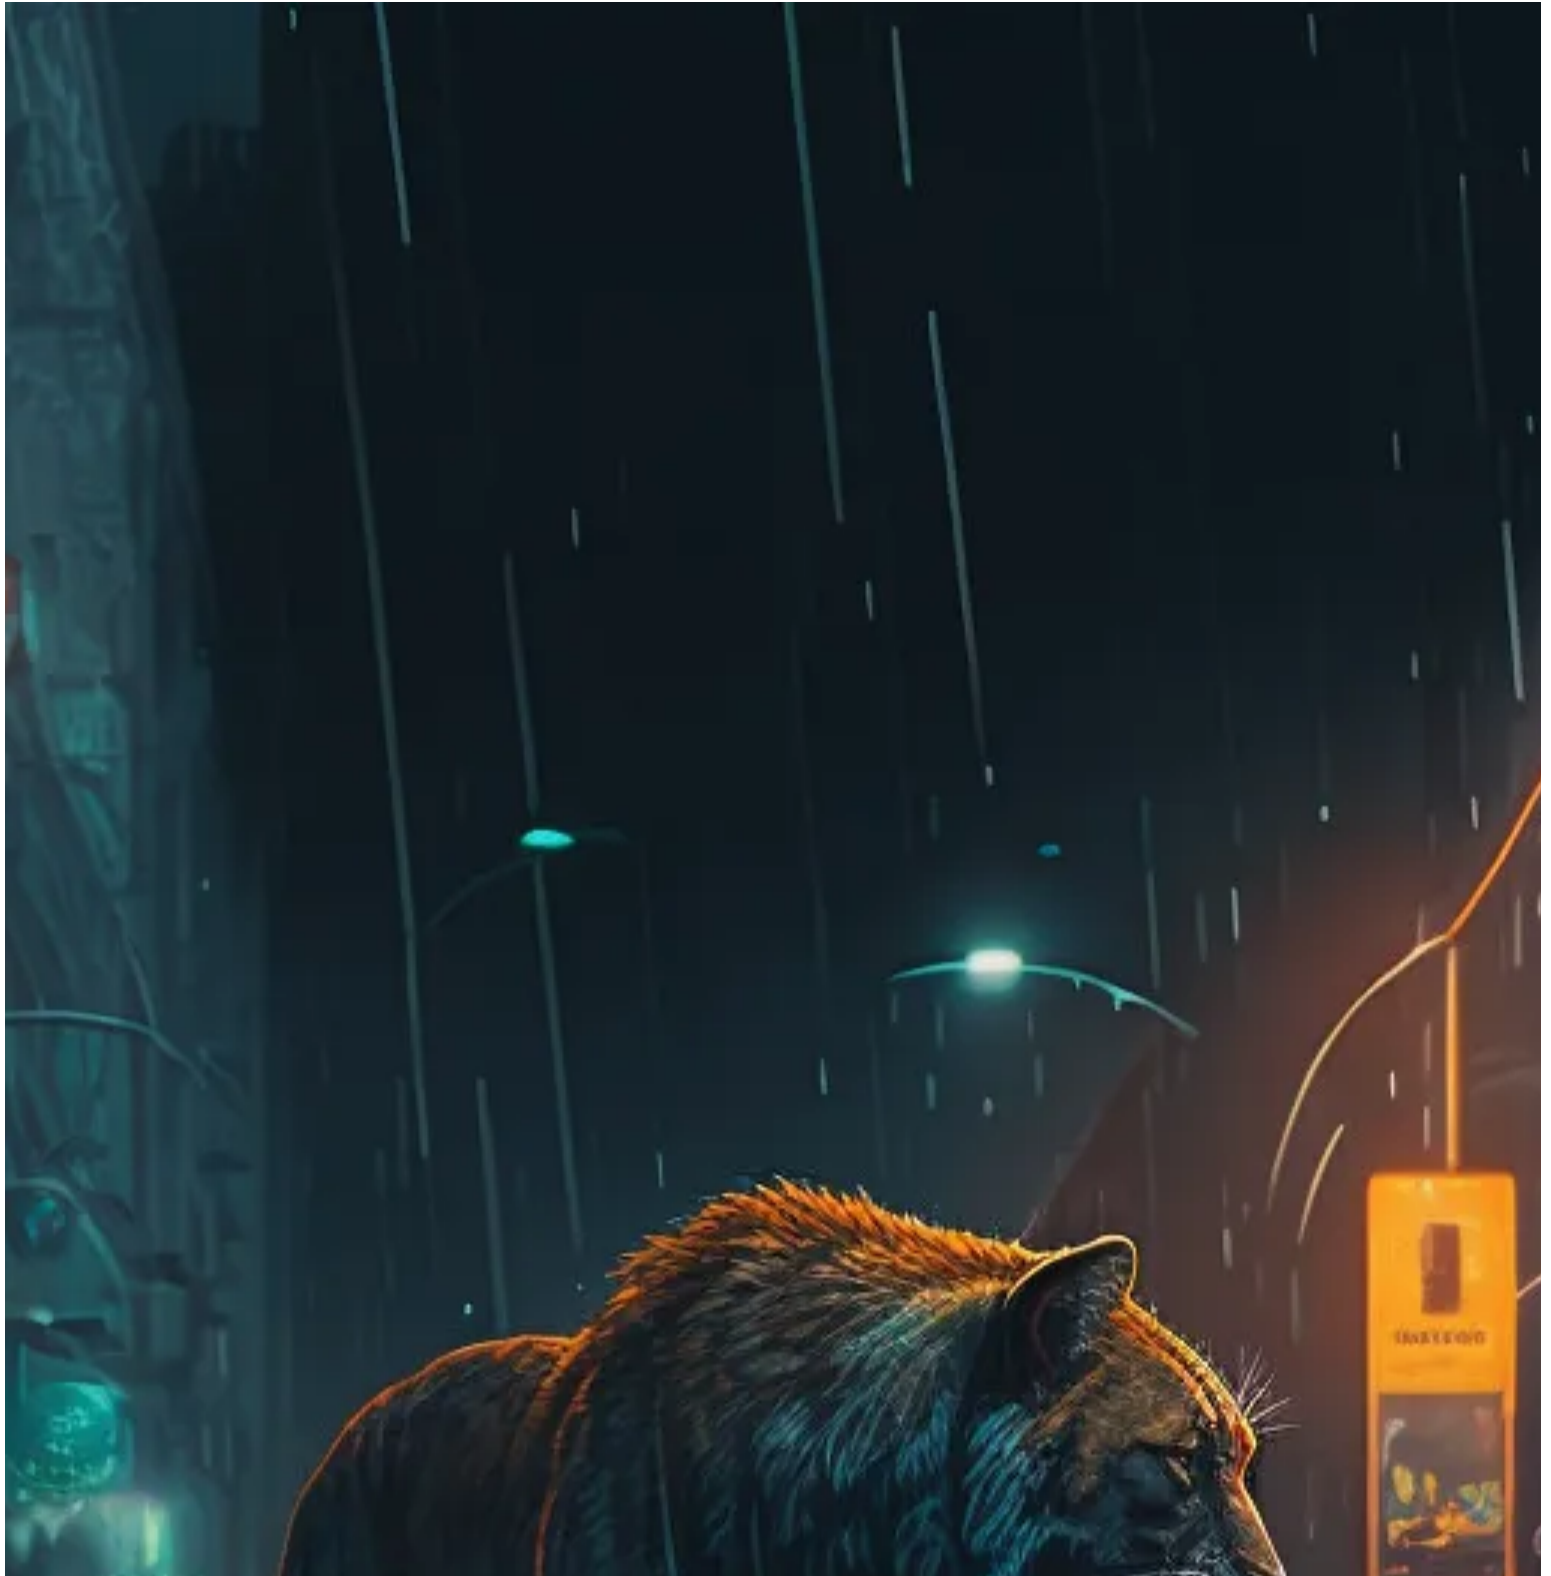
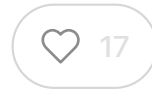
Glimpse of the Future: AI in Hollywood

The horse is out of the barn, and this is where it's heading



James Grimmelmann

Feb 7





JUST A MINUTE The author requested a “cyberpunk lion walking down a deserted urban streetscape in the rain at night,” and image-creating AI platform Midjourney output this near-instant rendering.

James Grimmelmann is the Tessler Family Professor of Digital and Information Law at Cornell Tech and Cornell Law School. He has degrees in computer science from Harvard and law from Yale. He teaches and writes about the Internet and intellectual property law.

This is the second part in our Glimpse of the Future series. Our first was about the “7

This is the second part in our *Glance of the Future* series. Our first was about the “Streamers” [from Ken Ziffren](#). The series is for paid subscribers.

If you go to [This Person Does Not Exist](#) and hit reload over and over, you will see an endless procession of humanity: confused kids with soup-bowl haircuts, middle-aged women with well-worn laugh lines, and fit younger guys smiling in the sunshine. But as the name of the site promises, none of these people are real. Every single face is fake. They were all produced by an algorithm that specializes in creating realistic human faces.

This Person Does Not Exist is a particular kind of AI called a “generative AI.” AI-driven apps have flourished in the last few years, ranging from the software in self-driving cars to [Merlin](#), which can identify any bird.

What makes generative AIs different is that they are creative. Merlin, for example, starts from a complex, messy input (a grainy photograph) and produces a simple output (“Grey Heron” for example).

But a generative AI runs in reverse: it starts with a simple input, called a “prompt,” and produces a rich media output. There are image AIs, music AIs, text AIs, programming AIs, and even video AIs. And they are remarkably good. If you feed the prompt “a Grey Heron” into OpenAI's image-generating [DALL-E](#), it outputs a perfectly convincing picture of a heron.

My own photo of a mystery bird ...





... was correctly identified by Merlin as a Grey Heron ...

Gray Heron



ADULT

Generally quite common and conspicuous in wetland habitats from marshes and tidal flats to small ponds, ditches, and wet fields; nests colonially in tall trees. Mainly seen as singles or in small groups, standing quietly in or at the edge of water, less often hunting...



... but when I typed "Grey Heron" into DALL-E ...





... it made a better photo of a Grey Heron than the one I took.

The most jaw-dropping generative AI at the moment is probably [ChatGPT](#), also from OpenAI. It can tell jokes, compose short stories and write rhyming instructions for assembling do-it-yourself furniture. But almost every area of human creativity is now being explored by generative AIs. In addition to DALL-E, image models like [Stable Diffusion](#) and [Midjourney](#) can produce remarkably beautiful art in hundreds of styles. [GitHub Copilot](#) can generate fluent, usable code in dozens of programming languages.

The pace of innovation in generative AIs is blistering. In just the last week, Google researchers released *two* groundbreaking musical AIs. One, called [MusicLM](#), can write and perform music from prompts like “slow tempo, bass-and-drums-led reggae song.” The other, [SingSong](#), takes an isolated vocal line and fills in a instrumental backing track. Meanwhile, [Drayk.it](#) will make a halfway-decent Drake rap about any subject you type in. Here's one about [Legos](#): “When I build I'm in control / My creativity is my only goal”). Listen here:

1x

0:00

-0:31



The entertainment industry has been roiled by one technological change after another: file-sharing, streaming, all-digital production. But generative AI may reshape every aspect of the business behind Hollywood, because it can affect every part of the creative process and upend every potential rights' dispute. A looming writers guild strike, ongoing fights over intellectual property, the prevalence of sequels and spinoffs

over original ideas — all of these could look very different if creatives are routinely using AI tools that can quickly spin a simple prompt out into a storyboard, a treatment, a temp track, a matte painting, or an effects shot. Can you feed an AI the scripts for the first 25 James Bond movies and ask it to write the next one? And if so, who owns the rights?

I don't know all the answers for what generative AI will mean for the entertainment business, but I can hopefully give some context for how to think through its implications. Below, I'll talk about:

- **What IP ownership looks like** in an AI-heavy world.
- How the roles of **writers and craftspeople** might shift and adapt.
- What AI might do to the **costs of production** and for the world of spinoffs and sequels.
- Real-life AI examples involving **Aaron Sorkin, *Fast and Furious*, *Seinfeld* and *Bridgerton*** that reveal just how much AI may change the way you work. And soon.

How Generative AI Works

The most successful generative AI models all use variations on the same basic technique, called “deep learning.” This type of model consists of a large network of individual nodes connected to each other, a bit like the neurons in a human brain. The model is “trained” by exposing it to a large number of example inputs, called “training data.” (For DALL-E, for example, each training example consists of an image and a caption describing it.) Each example strengthens the connections between the parts of the network that activated in response to it, like the connections between neurons strengthening each time they see a familiar face.

With a clever enough training strategy, a model can learn in a way that allows the process to be reversed. ChatGPT learns the patterns of human writing; its model represents the common ways that human-written passages go. This means that if a user prompts ChatGPT by giving it the start of a passage, like “It was a dark and

stormy night,” ChatGPT then, as you can see, continues the story:



It was a dark and stormy night



, the rain pounded against the windows as lightning lit up the sky. The wind howled through the trees, causing branches to creak and snap. It was a night unlike any other, a night of mystery and uncertainty.



For those of you in Hollywood, this is what happens when I ask ChatGPT to put its own spin on current events in the style of one of your business’ most highly-paid writers:



Write a comedic walk-and-talk scene in the style of Aaron Sorkin in which the president finds out that there is a Chinese high-altitude surveillance balloon floating over the United States.



INT. WHITE HOUSE - DAY

The President, JIM, walks quickly down the hall with his chief of staff, LUCY, by his side.

JIM: What do you mean there's a Chinese surveillance balloon floating over the United States?!

LUCY: (keeping pace with the president) That's exactly what I mean, sir. It was picked up by our radar systems this morning.

JIM: (stopping suddenly, turning to face Lucy) How could this happen? Don't we have a missile defense system in place to prevent this kind of thing?

LUCY: (urgently) Yes, we do, but this thing is high up, too high for our missiles to reach.

JIM: (sighs) Alright, let's get the Secretary of Defense in here.

LUCY: (pulls out her cell phone) On it, Mr. President.

JIM: (starts walking again) And what about the Chinese government, have we contacted them?

LUCY: (on the phone, walking and talking) We have, sir. They claim it's a weather balloon that got away from them.

JIM: (stops again, throwing up his hands) A weather balloon? That's the best they could come up with?

LUCY: (smirks) Apparently so, sir.

ChatGPT can also write pitches — although when I asked it for a *Fast and Furious* sequel, it got a little far from the series' roots:



Pitch the next Fast and Furious script.



Title: "Fast and Furious: Beyond the Speed"

In this high-octane adventure, our heroes are faced with their greatest challenge yet: traveling to the outer reaches of space to stop a rogue cyber-terrorist from destroying the planet.

With their fearless leader, Dominic Toretto, at the helm, the Fast and Furious crew must navigate through treacherous asteroid fields, outwit an army of hi-tech robots, and outmaneuver the cyber-terrorist's powerful spaceship.

As the race to save the world reaches new heights, our heroes discover that the only way to defeat their opponent is to push their vehicles to the limit and beyond. With a mix of heart-pumping action, high-tech gadgetry, and fierce loyalty, Fast and Furious: Beyond the Speed is a thrill ride that will have audiences on the edge of their seats.

How it Works

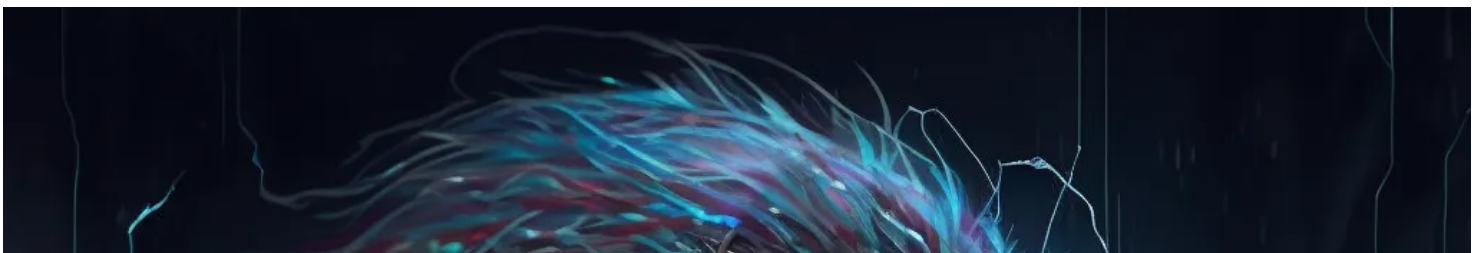
Successful generative AI models are enormous. OpenAI's GPT-3 model, on which ChatGPT is based, stores what it has learned using 175 billion different data points, called "parameters." AI models also require immense amounts of training data. Stable Diffusion was trained on 5.85 billion images scraped from the Internet.

Unsurprisingly, training is slow and expensive. The process can take months and cost millions of dollars. There is a race underway among major AI companies to scale up their models — rumors are that GPT-4 will have a trillion parameters. This is not a game that hobbyists and dilettantes can play.

But once a model has been trained, actually running it on a prompt to generate an output is cheap and fast. It takes less than a minute to generate an image using the most popular AIs, and the text models can respond in seconds. And while cutting-edge generative AIs like GPT-3 are kept closely guarded by their creators, there are plenty of consumer-grade generative AIs that are widely available, like the buzzy [Lensa Magic Avatar](#). Indeed, there are even plenty of smaller models that work perfectly well on a personal computer — or even a phone! Apple, for example, has designed [software and hardware](#) to make AI models run quickly on iPhones, supporting apps like the cute and friendly image maker [Draw Things](#).

What Generative AI Will Do

These models massively reduce the time, effort and cost required to produce content. It becomes trivially easy to have an AI generate text that could pass for something written by a human. The same goes for images — and very soon for audio and video.





FAST DELIVERY My first “cyberpunk lion” request. (Midjourney)

What is harder is to use AI to realize a specific creative vision. Anyone can type “cyberpunk lion” into DALL-E or Midjourney, but if you want the lion to be walking down a deserted urban streetscape in the rain at night, you need to tell that to the AI (and yes, it will do it, see top photo). There is already a large gap between AI amateurs (like myself) and professional prompt engineers who know how to write prompts that coax and tweak the AI into exactly generating what they have in mind. That gap will only grow.

This means that the coming flood of AI-generated content is not just one phenomenon. Different kinds of artists and authors will use generative AI in different

ways.

At the high end, many creative professionals will incorporate generative AIs into their creative process. Musicians and 3D artists already use powerful software tools like Logic Pro and Maya. Soon, generative AI will become part of the standard tooling creators are trained on and expected to know.

Indeed, existing tools will simply incorporate AI seamlessly: instead of moving a slider to adjust the color balance of an image manually in Adobe Photoshop, a digital artist will simply click a button to have AI make the image match the color palette of a reference photo — or have it extend a handful of trees in the background of a photo into a lush forest. Or, to generate a costume palette for a TV series.

I typed in “costume ideas for a *Bridgerton* Christmas party scene” and Midjourney generated this collage:





(Midjourney)

The VFX Factor

While many artists will foreswear AI tools (just as many painters today still work on a physical canvas rather than a tablet screen), in commercial settings, generative AI will be standard practice. Major movie and television productions will rely heavily on AI for their VFX work. Meanwhile, individuals and small teams will use generative AI to give shoestring productions the look and sound of productions with a thousand times their budget. (Something like this has already happened in video games. Triple-A titles like *God of War* have immense production teams and custom-built tech, but it is also possible for independent groups to make finely polished games using off-the-shelf tools like Unity.)

In the middle of the market, the constant struggle to capture consumers' attention will trap creators in an AI-powered arms race, where the goal is simply to crank out as much content as possible. It doesn't have to be good, just good enough. For a sense of this future, look at the YouTube treadmill, where influencers and streamers upload new videos as fast as they can.

Everyone whose job already involves writing under time pressure will feel the temptation of AI tools — or, if they do not, their bosses will. Advertising copywriters will ask AI to write website copy; journalists will ask AI to fill in the formulaic parts of stories. Romance novelists and travel bloggers will prompt an AI model with a basic plot sketch or a destination, skim the results, and then hit “publish.” And spammers and scammers will do the same, but won't even bother to slow down to look over the

results themselves. Now imagine the same for every creative medium. Right now, humans make replacement-level art. Very soon, AIs will be able to.

This coming firehose of AI-generated content will be shockingly good on the surface, and shockingly bad underneath. The news website CNET has been [secretly publishing](#) AI-“written” articles for months, some of which contained serious factual errors. CNET made corrections, but other content farms won't bother — or the errors will simply go undetected. (BuzzFeed, for one, is [full steam ahead](#) on creating content with generative AI.)

But not all generative AI will be directed at the marketplace. People have always made quirky personal art for themselves and their friends; AIs will help them be even quirkier. Parents of toddlers who love dinosaurs will ask ChatGPT for bedtime-story prompts. Couples will compose and perform Valentine's Day songs for each other — with a little help from AI songwriting and auto-tuning assistants. And fans everywhere will use AI tools to make their fanfic into actual crossover episodes of, say, long-cancelled series starring their favorite long-gone characters.



write fanfic about Mork from Mork and Mindy falling in love with Katniss from The Hunger Games



Mork from the planet Ork was a curious and playful being, always eager to learn about new cultures and ways of life. During his travels, he stumbled upon the world of Panem, where he met Katniss Everdeen, the brave and determined survivor of the Hunger Games.

At first, Mork was intrigued by Katniss's strength and courage, as well as her unwavering determination to fight against the oppressive regime of the Capitol. However, as he got to know her better, he realized that there was much more to her than just her fierce exterior.

Katniss had a kind and caring heart, and Mork couldn't help but fall in love with her. He admired her selflessness and her willingness to put the needs of others before her own.

Despite the obstacles they faced, including the cultural differences between their worlds, Mork and Katniss grew closer and closer, until their love for each other was unbreakable.

Together, they explored the world of Panem and beyond, facing new challenges and adventures along the way. And despite the odds, they proved that love knows no boundaries, and that it can conquer even the greatest of obstacles.

In the end, Mork and Katniss were able to build a life together, surrounded by the love of their friends and family, and the joy of their newfound happiness. And they knew that no matter where their adventures took them, they would always have each other by their side.

Is Generative AI Legal?

The copyright law of generative AI is unsettled. [Getty Images](#) and a [group of artists](#) have sued the creators of Stable Diffusion, whose training data included [photos scraped](#) from sites including Pinterest, Getty, Etsy, Adobe, and DeviantArt. Meanwhile, a [group of programmers](#) have sued OpenAI over GitHub Copilot, a generative AI that can write snippets of code and is already in wide use. But these lawsuits — and others that will follow — are likely to take years to resolve.


For the moment, AI companies have some powerful precedents to rely on. In [2014](#) and [2015](#), a federal appeals court held that it was fair use for Google to scan millions of books to create its [Google Books search engine](#). Although Google made complete digital copies of books without the authors' or publishers' permission, the search engine broke each book up into “discontinuous, tiny fragments” that “communicate[d] little of the sense of the original.” Another appeals court reached a similar result in [2009](#), allowing a plagiarism-detecting company to build a database of high-school term papers that teachers could check student work against.

There is a strong argument that training an AI model is legal, even when copyrighted works are part of the training data. Like building a search engine or a plagiarism

works are part of the training data. Like building a search engine or a plagiarism checker, it is a “transformative” use. These algorithms are not “experiencing” the works they train on, reading history books to understand the origins of the civil rights movement or looking at photographs to admire their composition. Instead, they are “analyzing” their training data to draw conclusions about that data — much like the kinds of scientific and educational uses that copyright law has always favored.

But even if “training” an AI model is fair use, using that model to “generate” works may not be. (This is the crucial difference between modern generative AIs and these earlier cases — a program that can “write” a term paper is very different from a program that can “read” one.) People have used generative AIs to produce photographs in the style of Annie Liebovitz and [verbatim excerpts from published textbooks](#). There is something about the style of *Seinfeld* that AIs seem to be particularly good at imitating. [Nothing, Forever](#) is an infinite livestream of an endless AI-generated *Seinfeld* episode, while other users have figured out that anything is funny as a *Seinfeld* scene, even programming. It’s clear that AIs can produce outputs that are highly similar to existing copyrighted works.



 **Riley Goodside**
@goodside

OpenAI's new ChatGPT writes a Seinfeld scene in which Jerry needs to learn the bubble sort algorithm:

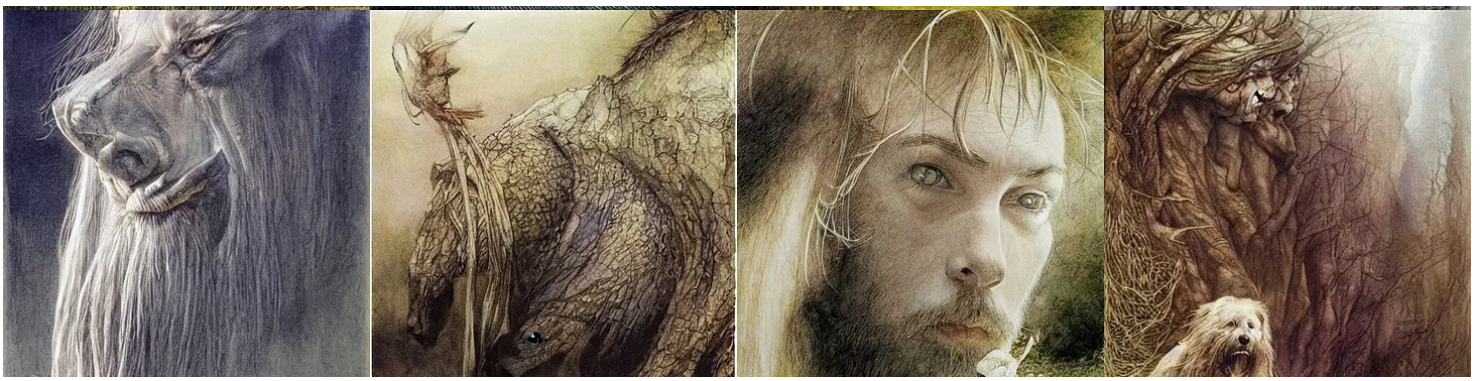
Write the complete script of a Seinfeld scene in which Jerry needs to learn the bubble sort algorithm. Include frequent jokes and note audience laughter in the script.	Elaine: (laughs) Bubble sort? That's a piece of cake. You just compare each element to the one next to it and swap them if they're in the wrong order.
INT. MONK'S CAFE - DAY	(audience laughs)
Jerry is sitting at the counter with George.	Jerry: (frustrated) Compare and swap? This is starting to sound like rocket science.
Jerry: (sighs) I can't believe I have to learn the bubble sort algorithm for my computer science class.	The waiter approaches the group.
George: (laughs) Bubble sort? That's the most basic sorting algorithm there is. Even a monkey could do it.	Waiter: Can I get you guys anything?
(audience laughs)	Jerry: (sighs) Just a coffee. I need to stay awake for this bubble sort lesson.
Jerry: Yeah, well I'm not a monkey, I'm a comedian. (picks up menu)	(audience laughs)
Kramer enters and sits down next to Jerry.	George: I'll have the same.
Kramer: Hey, what's so funny?	Jerry: (to Elaine) Hey, maybe you can help me with the bubble sort. You seem to know what you're talking about.
Jerry: I have to learn the bubble sort algorithm for my computer science class.	Elaine: (smirks) Sorry Jerry, I have better things to do than teach you a basic sorting algorithm.
Kramer: (laughs) The bubble sort? That's easy. You just keep swapping adjacent elements until the list is sorted.	(audience laughter)
(audience laughs)	Jerry: (deflated) Fine, I'll just figure it out on my own.
	The group continues their conversation as the waiter brings their food and drinks.
	(audience laughter)

10:11 PM · Nov 30, 2022

4,192 Likes 472 Retweets

Where do these similarities come from? In some cases, the model has simply memorized some of the training data: give it the right prompt and it regurgitates the original work, word for word or bit for bit. In other cases, the model has learned to recognize a distinctive artistic style; some of those styles are broad categories like “sports photography” but others are highly specific, like [Alan Lee's Lord of the Rings art](#). And in still other cases, it is the “prompt” that contains the kernel of infringement: give DALL-E an image to start with and it can “[outpaint](#)” to extend the scene beyond its original boundaries.





LORD OF THE RINGS Stable Diffusion outputs in the style of Alan Lee.

[\(source\)](#)

It is unlikely that courts will give these generations a free pass just because they were made with an AI rather than by hand. Of course, many of them will still be fair uses; a kid can put their own face on Iron Man with scissors and glue, with cut and paste in Photoshop, or with a generative AI. But there is likely to be a core of uses that are decidedly not fair, like using an AI to design a line of beer cans in the style of **Ralph Steadman** that competes with the ones he illustrated for [Flying Dog Brewery](#). And depending on what the Supreme Court does in its pending [Warhol v. Goldsmith](#) case (about a **Lynn Goldsmith** photo of **Prince** that **Andy Warhol** used as the starting point for 16 silkscreen prints), the category of infringing uses may be quite large.

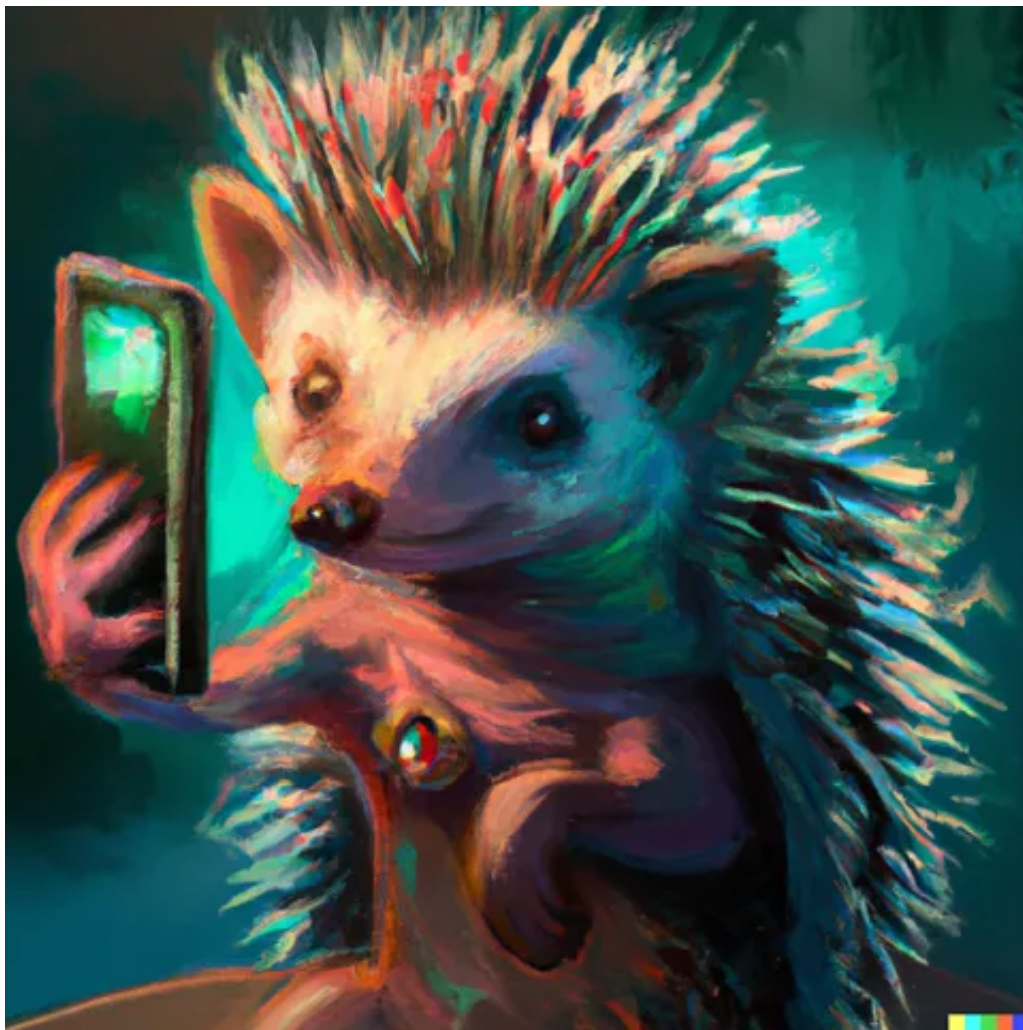
This split — legal to train, but not necessarily to use — may create massive AI risks. For one thing, copyright owners will ask courts to put restrictions on how AI models can be released. It is possible, for example, that a court could order OpenAI to add filters preventing users from generating [Dr. Seuss-style rhyming poems](#)

This is a bit like what happened to Napster. When Napster was [ordered](#) to remove the record companies' songs, it tried, failed, and shut down. Here, no one, not even the technical experts at AI companies, fully understands exactly how their models work, or “where” in a model the training data is represented. Fully effective copyright filtering is currently impossible.

Even if it were possible to filter an AI’s outputs, several powerful models (most prominently Stable Diffusion) have already been released publicly as open source. Anyone can download them and run these models on their own device! Similarly, anyone can tinker with the models to make them more powerful or [disable their](#)

anyone can tinker with the models to make them more powerful, or [disable their safety filters](#). Just like with file-sharing, the horse is out of the barn. No matter how the models were trained, and no matter what the courts do, copyright owners will not be able to shut down generative AI entirely.

AI companies would prefer that the responsibility for avoiding infringement rest with users. But the black-boxiness of generative AIs means that using one is always a little bit at your own risk. Maybe the oil painting of a hedgehog taking a selfie that I generated using DALL-E is unique, or maybe it's a near-exact ripoff of someone else's oil painting of a hedgehog taking a selfie.



ORIGINAL AI? No idea if my DALL-E request for an oil painting hedgehog taking a selfie is unique or a ripoff of an image scraped from the internet.

Many organizations already have policies on how they use open-source software to avoid legal risks. They will need to develop similar policies on the use of generative

AIs, telling employees whether, when, and how they can incorporate AI-produced content into their workflows. And then they will have the extraordinarily difficult challenge of enforcing those policies for employees who have generative AI tools literally at their fingertips and are accustomed to using them constantly in their everyday lives.

Final Words

Generative AI may not transform human creativity. So far, at least, people are mostly using AI tools to make familiar kinds of art, and just more of it. But AI is likely to permeate the human social world, just as the Internet and social media have done. They are always with us, never far away, and all but impossible to avoid.

3 Comments



1 more comment...
