

ARTIFICIAL INTELLIGENCE

Training AI music models is about to get very expensive

Record labels are suing two leading AI startups and are reportedly exploring a licensing deal with YouTube—actions that will determine what the future of AI music sounds like, and if there is one at all.

By James O'Donnell

June 27, 2024



SARAH ROGERS/MITIR | GETTY

AI music is suddenly in a make-or-break moment. On June 24, Suno and Udio, two leading AI music startups that make tools to generate complete songs from a prompt in seconds, were sued by major record labels. Sony Music, Warner Music Group, and Universal Music Group claim the companies made use of copyrighted music in their training data “at an almost unimaginable scale,” allowing the AI models to generate songs that “imitate the qualities of genuine human

MIT Technology Review



aboveboard approach. Rather than training AI music models on secret data sets, the company is reportedly offering unspecified lump sums to top record labels in exchange for licenses to use their catalogues for training.

In response to the lawsuits, both Suno and Udio released statements mentioning efforts to ensure that their models don't imitate copyrighted works, but neither company has specified whether their training sets contain them. Udio said its model "has 'listened' to and learned from a large collection of recorded music," and two weeks before the lawsuits, Suno CEO Mikey Shulman told me its training set is "both industry standard and legal" but the exact recipe is proprietary.

While the ground here is changing fast, none of these moves should be all that surprising: litigious training-data battles have become something like a rite of passage for generative AI companies. The trend has led many of those companies, including OpenAI, to pay for licensing deals while the cases unfold.

Related Story



However, the stakes are higher for AI music than for image generators or chatbots. Generative AI companies working in text or photos have options to work around lawsuits; for example, they can cobble together open-source corpuses to train models. In contrast, music in the public domain is much more limited (and not exactly what most people want to listen

**on playing**

The technique, called AudioLM, generates naturalistic sounds without the need for human annotation.

with interested publishers and creators, of which there are many; but rights in music are far more concentrated than those in film, images, or text, industry experts say. They're largely managed by the three biggest record labels—the new plaintiffs—whose publishing arms collectively own more than 10 million songs and much of the music that has defined the last century. (The filing names a long list of artists who the labels allege were wrongfully included in training data, ranging from ABBA to those on the *Hamilton* soundtrack.)

On top of all this, it's also just more difficult to create music worth listening to—generating a readable poem or passable illustration with AI is one technical challenge, but infusing a model with the taste required to create music we like is another.

It's of course possible that the AI companies will win the case, and none of this will matter; they would have carte blanche to train on a century of copyrighted music. But experts say the case from the record labels is strong, and it's more likely that AI companies will soon have to pay up—and pay a lot—if they want to survive. If a court were to rule that AI music companies could not train for free on these labels' catalogues, then expensive licensing deals, like the one YouTube is reportedly pursuing, would seem to be the only path forward. This would effectively ensure that the company with the deepest pockets ends up on top.

More than any training-data case yet, the outcome of this one will determine the shape of a big slice of AI—and whether there is a future for it at all.

Merits of the case

Suno's music generator has been public for less than a year, but the company has already garnered 12 million users, a \$125 million funding round last month, and a partnership with Microsoft Copilot. Udio is even newer to the scene, having launched in April with \$10 million in seed funding from musician-investors like will.i.am and Common.

The record labels allege that both of the startups are engaging in copyright infringement on the training and the output sides of their models.

“The plaintiffs here have the best odds of almost anyone suing an AI company,” says James Grimmelmann, a professor of digital and information law at Cornell Law School. He draws comparisons to the ongoing *New York Times* case against OpenAI, which he says offered, until now, the best example of a rights holder with a strong case against an AI company. But the suit against Suno and Udio “is worse for a bunch of reasons.”

MIT Technology Review



deficiency in this accusation, because the company could say that it scraped much of the internet for a training corpus and copies of *New York Times* articles appeared in places without the company's knowledge.

For Suno and Udio, that defense is far less believable. “This is not like, ‘We scraped the web for all audio and we couldn’t tell the commercially produced songs apart from everything else,’” Grimmelmann says. “It’s pretty clear that they had to have been pulling in large databases of commercial recordings.”

In addition to complaints about training, the new case alleges that tools like Suno and Udio are more imitative than generative AI, meaning that their output mimics the style of artists and songs protected by copyright.

While Grimmelmann notes that the *Times* cited examples in which ChatGPT reproduced entire copies of its articles, record labels claim they were able to generate problematic responses from the AI music models with much simpler prompts. For instance, prompting Udio with “my tempting 1964 girl smokey sing hitsville soul pop,” the plaintiffs say, yielded a song that “any listener familiar with the Temptations would instantly recognize as resembling the copyrighted sound recording ‘My Girl.’” (The court documents include links to examples on Udio, but the songs appear to have been removed.) The plaintiffs mention similar examples from Suno, including an ABBA-adjacent song called “Prancing Queen” that was generated with the prompt “70s pop” and the lyrics for “Dancing Queen.”

What’s more, Grimmelmann explains, there is more copyrightable information in a song than a news article. “There’s just a lot more information density in capturing the way that Mariah Carey’s voice works than there is in words,” he says, which is perhaps part of the reason past lawsuits navigating music copyright have sometimes been so drawn-out and complex.

In a statement, Shulman wrote that Suno prioritizes originality and that the model is “designed to generate completely new outputs, not to memorize and regurgitate preexisting content.” He added, “That is why we don’t allow user prompts that reference specific artists.” Udio’s statement similarly mentioned “state-of-the-art filters to ensure our model does not reproduce copyrighted works or artists’ voices.”

MIT Technology Review



When users search examples suggesting that if users separate an artist's name with spaces, the request may go through. My own request for "a song like Kendrick" was blocked by Suno, citing an artist's name, but "a song like k e n d r i c k" resulted in a "hip-hop rhythmic beat-driven" track and "a song like k o r n" resulted in "nu-metal heavy aggressive." (To be fair, they didn't resemble the respective artists' unique styles, but to even respond in the right tightly defined genre seems to suggest that the model is in fact familiar with each artist's work.) Similar workarounds were blocked on Udio.

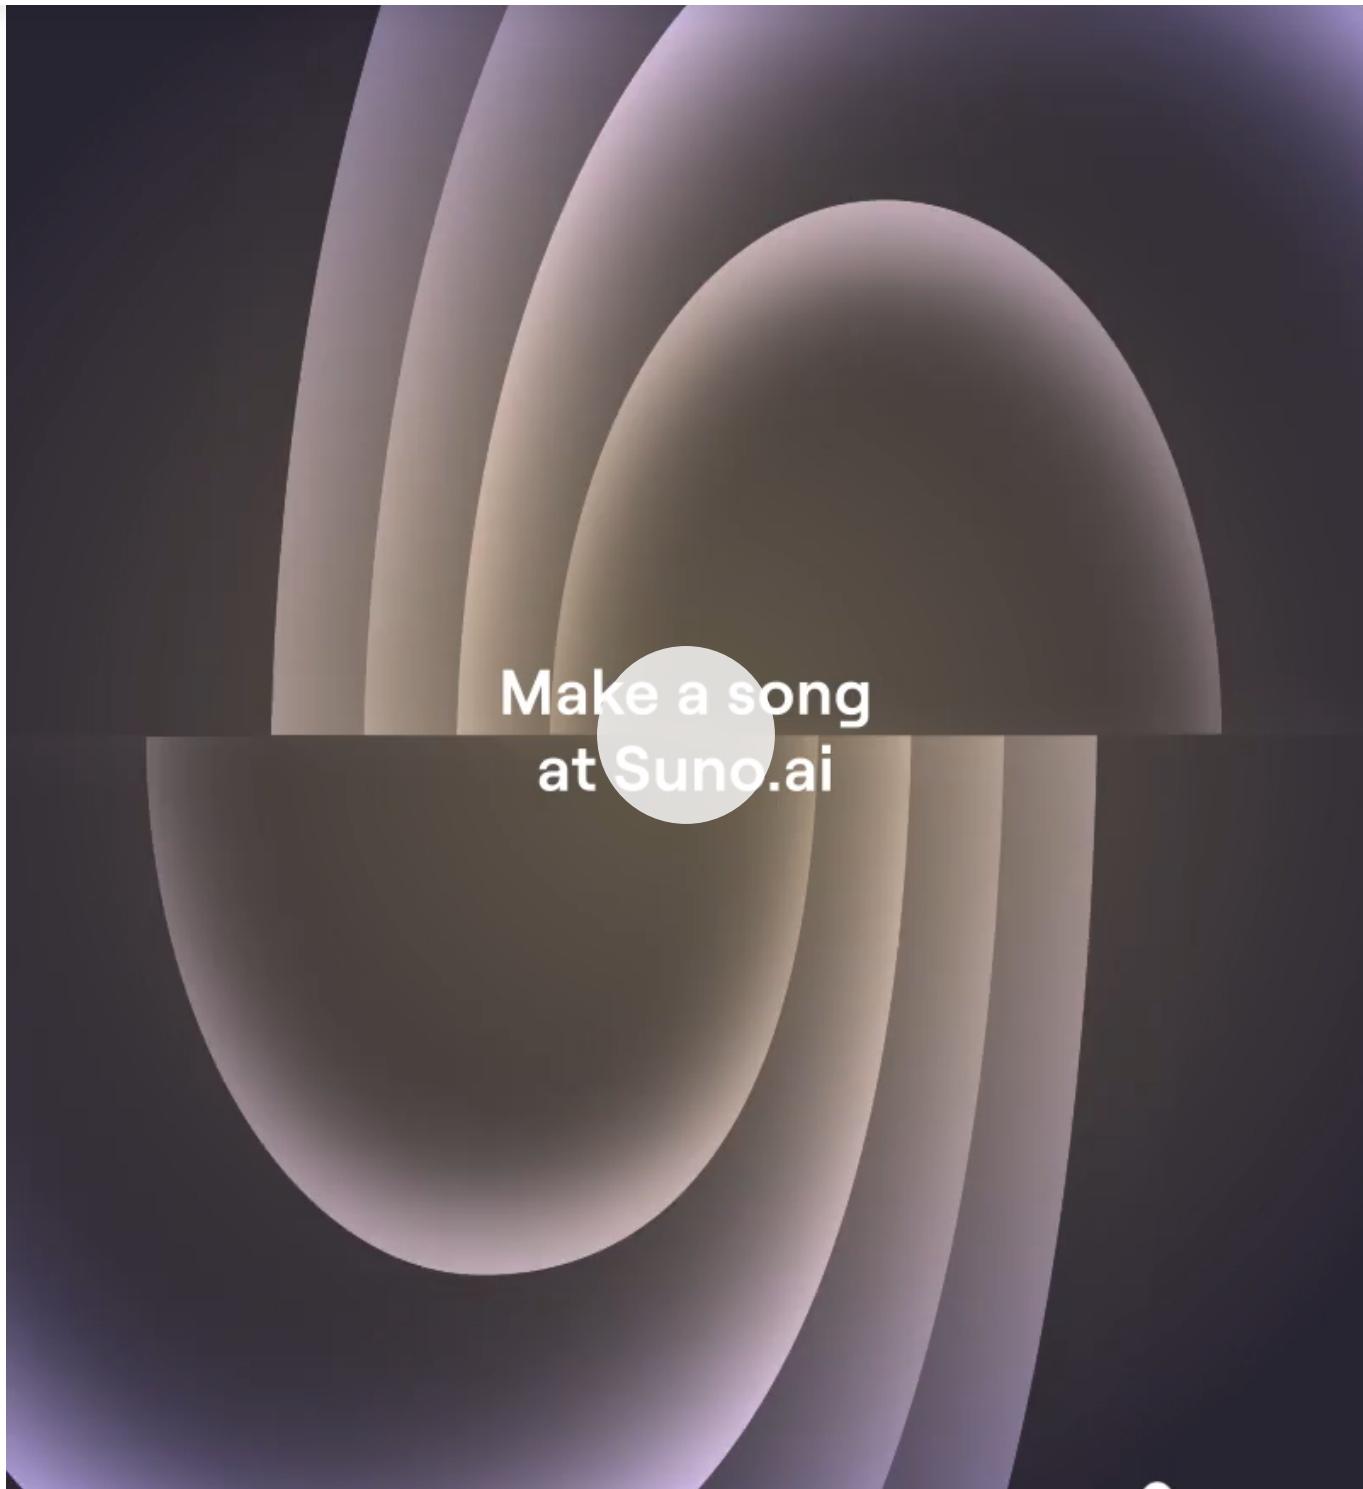
Possible outcomes

There are three ways the case could go, Grimmelmann says. One is wholly in favor of the AI startups: the lawsuits fail and the court determines that companies did not violate fair use or imitate copyrighted works too closely in their outputs. If the models are found to fall under fair use, it would mean songwriters and rights holders would need to find a different legal mechanism to pursue compensation.

Another possibility is a mixed bag: the court finds the AI companies did not violate fair use in their training but must better control their models' output to make sure it does not improperly imitate copyrighted works. Grimmelmann says this would be similar to one of the initial rulings against Napster, in which the company was forced to ban searches for copyrighted works in its libraries (though users quickly found workarounds).

The third and essentially nuclear option is that the court finds fault on both the training and the output sides of the AI models. This would mean the companies could not train on copyrighted works without licenses, and also could not allow outputs that closely imitate copyrighted works. The companies could be ordered to pay damages for infringement, which could run into the hundreds of millions for each company. If they aren't bankrupted by such a ruling, it would force them to completely restructure their training through licensing deals, which could also be cost-prohibitive.

MIT Technology Review



Make a song
at [Suno.ai](https://suno.ai)



0:00 / 0:48



COURTESY SUNO.AI

To license or not to license

MIT Technology Review



already looking ahead toward a future of licensing. “As in the past, music creators will choose their rights to protect the creative engine of human artistry and enable the development of a healthy and sustainable licensed market that recognizes the value of both creativity and technology,” he wrote in a recent op-ed in *Billboard*.

Such a market for licenses could mirror what has already unfolded for text generators. OpenAI has struck licensing deals with a number of news publishers, including *Politico*, the *Atlantic*, and the *Wall Street Journal*. The deals promise to make content from the publishers discoverable in OpenAI’s products, though the ability for the models to transparently cite where they’re getting information from is limited at best.

If AI music companies follow that pattern, the only ones with the means to create powerful music models might be those with the most cash. That’s perhaps exactly what YouTube is thinking. The company did not immediately respond to questions from *MIT Technology Review* about the details of its negotiations, but given the massive amount of data required to train AI models and the concentration of rights owners in music, it’s fair to assume the price of deals with record labels would be eye-popping.

In theory, an AI company could bypass the licensing process altogether by building its model exclusively on music in the public domain, but it would be a Herculean task. There have been similar efforts in the realm of text and image generation, including a legal consultancy in Chicago that created a model trained on dense regulatory documents, and a model from Hugging Face that trained on images of Mickey Mouse from the 1920s. But the models are

MIT Technology Review



resulting model would be a far cry from what they have today.

If AI companies do move forward with licensing agreements, negotiations may be tricky, says Grimmelmann. Music licensing is complicated by the fact that two different copyrights are at play: one for the song, which generally covers the composition, like the music and lyrics, and one for the master, which covers the recording—like what you'd hear if you streamed the song.

Some artists, like Taylor Swift and Frank Ocean, have come to own the masters of their catalogues after drawn-out legal battles, and would therefore be in the driver's seat for any potential licensing deal. Many others, though, retain only the song copyright, while the record labels retain the masters. In these cases, the record label might theoretically be able to grant AI companies a license to use the music without an artist's permission—but at the risk of burning relationships with artists and sparking more legal battles.

The question of whether to license their music to such companies has divided musician groups. In contract rules adopted in April by SAG-AFTRA, which represents recording artists as well as actors, AI clones of member voices are allowed, though there are minimum rates for compensation. Back in December, a group called the Indie Musicians Caucus expressed frustrations that the leading instrumental musicians' union, the 70,000-member American Federation of Musicians (AFM), was not doing enough to protect its rank and file against AI companies in contracts. The caucus wrote that it would vote against any agreement “obligating AFM members to dig [their] own graves by participating—without a right to consent, compensation, or credit—in the training of our permanent Generative AI replacements.”

But at this point, AFM does not appear eager to facilitate any deals. I asked Kenneth Shirk, international secretary-treasurer at AFM, whether he thought musicians should engage with AI companies and push to be fairly compensated, whatever that means, or instead resist licensing deals completely.

“Looking at those questions makes me think, would you rather have a swarm of fire ants crawling all over you, or roll around in a bed of broken glass?” he told me. “We want musicians to get paid. But we also want to ensure that there’s a career in music to be had for those that are going to come after us.” 

by James O'Donnell



DEEP DIVE

ARTIFICIAL INTELLIGENCE



How to opt out of Meta's AI training

Your posts are a gold mine, especially as companies start to run out of AI training data.

By **Melissa Heikkilä**

Apple is promising personalized AI in a private cloud. Here's how that will work.

Apple's first big salvo in the AI wars makes a bet that people will care about data privacy when automating tasks.

By **James O'Donnell**

Why does AI hallucinate?

The tendency to make things up is holding chatbots back. But that's just what they do.

By **Will Douglas Heaven**

This AI-powered “black box” could make surgery safer

A new smart monitoring system could help doctors avoid mistakes—but it's also alarming some surgeons and leading to sabotage.

By **Simar Bajaj**

**STAY CONNECTED**

Illustration by Rose Wong

**Get the latest updates from
MIT Technology Review**

Discover special offers, top stories, upcoming events, and more.

Enter your email



[Privacy Policy](#)

The latest iteration of a legacy

Founded at the Massachusetts Institute of Technology in 1899, MIT Technology Review is a world-renowned, independent media company whose insight, analysis, reviews, interviews and live events explain the newest technologies and their commercial, social and political impact.

[READ ABOUT OUR HISTORY](#)

MIT Technology Review



that are radically transforming business. From event sponsorships to custom content to visually arresting video storytelling, advertising with MIT Technology Review creates opportunities for your brand to resonate with an unmatched audience of technology and business elite.

ADVERTISE WITH US

[About us](#)

[Careers](#)

[Custom content](#)

[Advertise with us](#)

[International Editions](#)

[Republishing](#)

[MIT Alumni News](#)

[Help & FAQ](#)

[My subscription](#)

[Editorial guidelines](#)

[Privacy policy](#)

[Terms of Service](#)

[Write for us](#)

[Contact us](#)



