

# OpenAI and journalism

We support journalism, partner with news organizations, and believe The New York Times lawsuit is without merit.

January 8, 2024

## [Safety & Alignment](#)

Our goal is to develop AI tools that [empower people](#) to solve problems that are otherwise out of reach. People worldwide are already using our technology to [improve their daily lives](#). Millions of developers and more than 92% of Fortune 500 are building on our products today.

While we disagree with the claims in The New York Times lawsuit, we view it as an opportunity to clarify our business, our intent, and how we build our technology. Our position can be summed up in these four points, which we flesh out below:

1. We collaborate with news organizations and are creating new opportunities
2. Training is fair use, but we provide an opt-out because it's the right thing to do
3. "Regurgitation" is a rare bug that we are working to drive to zero
4. The New York Times is not telling the full story

## **1. We collaborate with news organizations and are creating new opportunities**

We work hard in our technology design process to support news organizations. We've met with dozens, as well as leading industry organizations like the News/Media Alliance, to explore opportunities, discuss their concerns, and provide solutions. We aim to learn, educate, listen to feedback, and adapt.

Our goals are to support a healthy news ecosystem, be a good partner, and create mutually beneficial opportunities. With this in mind, we have pursued partnerships with news organizations to achieve these objectives:

1. Deploy our products to benefit and support reporters and editors, by assisting with time-consuming tasks like analyzing voluminous public records and translating stories.
2. Teach our AI models about the world by training on additional historical, non-publicly available content.
3. Display real-time content with attribution in ChatGPT, providing new ways for news publishers to connect with readers.

Our early partnerships with the [Associated Press](#), [Axel Springer](#), [American Journalism Project](#) and [NYU](#) offer a glimpse into our approach.

## **2. Training is fair use, but we provide an opt-out because it's the right thing to do**

Training AI models using publicly available internet materials is fair use, as supported by long-standing and widely accepted precedents. We view this principle as fair to creators, necessary for innovators, and critical for US competitiveness.

The principle that training AI models is permitted as a fair use is supported by a wide range of [academics](#), [library associations](#), [civil society groups](#), [startups](#), [leading US companies](#), [creators](#), [authors](#), and [others](#) that recently submitted comments to the US Copyright Office. Other regions and countries, including the [European Union](#), [Japan](#), [Singapore](#), and [Israel](#) also have laws that permit training models on copyrighted content—an advantage for AI innovation, advancement, and investment.

That being said, legal right is less important to us than being good citizens. We have led the AI industry in providing a simple opt-out [process](#) for publishers (which The New York Times adopted in August 2023) to prevent our tools from accessing their sites.

### **3. “Regurgitation” is a rare bug that we are working to drive to zero**

Our models were designed and trained to learn concepts in order to apply them to [new problems](#).

Memorization is a rare failure of the learning process that we are continually making progress on, but it’s more common when particular content appears more than once in training data, like if pieces of it appear on lots of different public websites. So we have measures in place to limit inadvertent memorization and prevent regurgitation in model outputs. We also expect our users to act responsibly; intentionally manipulating our models to regurgitate is not an appropriate use of our technology and is against our terms of use.

Just as humans obtain a broad education to learn how to solve new problems, we want our AI models to observe the range of the world’s information, including from every language, culture, and industry. Because models learn from the enormous aggregate of human knowledge, any one sector—including news—is a tiny slice of overall training data, and any single data source—including The New York Times—is not significant for the model’s intended learning.

### **4. The New York Times is not telling the full story**

Our discussions with The New York Times had appeared to be progressing constructively through our last communication on December 19. The negotiations focused on a high-value partnership around real-time display with attribution in ChatGPT, in which The New York Times would gain a new way to connect with their existing and new readers, and our users would gain access to their reporting. We had explained to The New York Times that, like any single source, their content didn't meaningfully contribute to the training of our existing models and also wouldn't be sufficiently impactful for future training. Their lawsuit on December 27—which we learned about by reading The New York Times—came as a surprise and disappointment to us.

Along the way, they had mentioned seeing some regurgitation of their content but repeatedly refused to share any examples, despite our commitment to investigate and fix any issues. We’ve demonstrated how seriously we treat this as a priority, such as in July when we [took down a ChatGPT feature](#) immediately after we learned it could reproduce real-time content in unintended ways.

Interestingly, the regurgitations The New York Times induced appear to be from years-old articles that have proliferated on [multiple third-party websites](#). It seems they intentionally manipulated prompts, often including lengthy excerpts of articles, in order to get our model to regurgitate. Even when using such prompts, our models don’t typically behave the way The New York Times

insinuates, which suggests they either instructed the model to regurgitate or cherry-picked their examples from many attempts.

Despite their claims, this misuse is not typical or allowed user activity, and is not a substitute for The New York Times. Regardless, we are continually making our systems more resistant to adversarial attacks to regurgitate training data, and have already made much progress in our recent models.

We regard The New York Times' lawsuit to be without merit. Still, we are hopeful for a constructive partnership with The New York Times and respect its long history, which includes reporting the [first working neural network](#) over 60 years ago and championing First Amendment freedoms.

We look forward to continued collaboration with news organizations, helping elevate their ability to produce quality journalism by realizing the transformative potential of AI.