The Hypocrisy at the Heart of the AI Industry

Alex Reisner

Editor's note: This work is part of AI Watchdog, The Atlantic's ongoing investigation into the generative-AI industry.

In April 2024, Eric Schmidt, the former Google CEO and a current AI evangelist, gave a closed-door lecture to a group of Stanford students. If these young people hoped to be Silicon Valley entrepreneurs, Schmidt explained, then they should be prepared to breach some ethical boundaries.

At that point, 19 lawsuits had been filed against generative-AI companies for copyright infringement, alleging that Anthropic, OpenAI, and others had stolen books and other media to train their generative models. Yet Schmidt told the students to go ahead and download whatever they need to build an accurate "test" version of their AI product. If the product takes off, "then you hire a whole bunch of lawyers to go clean the mess up," he said. "If nobody uses your product, then it doesn't matter that you stole all the content."

Stanford posted a video of the talk on YouTube in August 2024, but it was removed a day later. (Stanford did not respond to my request for comment about the removal.)

When I recently obtained a copy, I was struck by Schmidt's readiness to say the quiet part out loud. He was articulating an attitude that is common in Silicon Valley but is usually stated as a legal or philosophical argument. When I reached one of Schmidt's spokespeople, they defended his position by telling me that Schmidt believes that the "fair use" of copyrighted work drives innovation. Others in the industry have cited the techno-libertarian idea that "information wants to be free," a frequently misunderstood credo that portrays information as a natural resource that should flow without restriction to whoever can use it.

But the credo never seems to apply to Silicon Valley's own information, whether it's the troves of personal data that companies have collected about us or the software they write. Photoshop, for example, doesn't want to be free. In fact, Photoshop is one of thousands of tech-industry products that are protected by patents. Inventions such as Google's original search algorithm and even design details, such as the "rounded rectangle" shape of Apple's iPhone, have also been patented, and companies employ teams of high-end attorneys to prosecute infringements.

The industry has long been a kind of intellectual-property battle zone, where damages in lawsuits frequently exceed nine figures. In 2017, for example, Waymo, Google's self-driving-car company, alleged that a former employee had stolen "confidential files and trade secrets, including blueprints, design files and testing documentation" for self-driving cars that were eventually shared with Uber. The case was settled for roughly $245 million. In the 2010s, Apple sued Samsung for copying elements of the iPhone and was initially awarded more than $1 billion in a patent-infringement battle that lasted seven years. Apple and Qualcomm have sued each other over IP in so many jurisdictions that it's hard to track.

In the pursuit of generative AI, tech companies have recently turned their aggressive strategies toward less prepared industries. As my reporting has shown, many top AI models have been trained on data sets containing massive numbers of copyrighted books, videos, and other works. This large-scale piracy has been excused in a number of ways: OpenAI (which has a corporate partnership with *The Atlantic*'s business team) has claimed that the company uses "publicly available information" to train its models; Anthropic has said that it has used books, but not in any commercial products; and Meta admits that it has used books in commercial products, but that doing so was "quintessential fair use."

Even as they claim the right to train their models on work belonging to other people, the AI companies have rejected similar reasoning when it comes to their own products. Consider OpenAI's terms of service for ChatGPT, which forbid use of the bot's "output to develop models that compete with OpenAI." Anthropic, Google, and xAI have similar clauses forbidding people from using the material generated by their chatbots to train competing products. In other words: We can train on your work, but you can't train on ours.

In the current economic environment, it's not surprising that companies vying for market dominance would operate with standards that serve their bottom line. But it's striking nonetheless how sharply their actions can contradict their professed values. Meta apparently does not want copies of its models on the web, even though it claims those models are "open," a word that typically means software is free and publicly available, and that implies a degree of goodwill or generosity on the part of the creator. It has reportedly sent notices demanding the deletion of such copies from online platforms. (Meta did not respond to a request for comment.)

Companies also know the value of training data, and at least one of them foresaw the backlash that taking such data might create. In 2021, one year before OpenAI released ChatGPT and two years before my reporting first revealed what was being used as AI-training data, Anthropic CEO Dario Amodei wrote an internal memo titled "An Economic Model for Compensating Data Producers." (It was recently unsealed in a copyright-infringement lawsuit against the company.) In the document, Amodei acknowledges that AI could be "an increasingly extractive concentrator of wealth" and that creators might eventually "grumble" or "get mad" as this fact becomes apparent. Resistance from creators might slow down AI progress, Amodei writes, and for this reason, he suggests compensating them "with a fraction of the profits from the model produced." Giving creators equity in the company could be a "great fit" for Anthropic's "public benefit orientation," Amodei wrote. Today, Anthropic still claims to provide a public benefit, but it has argued in court that using copyrighted books is "fair use"—meaning, essentially, that the authors are entitled to nothing. Anthropic declined to comment when I reached out for this article.

Companies argue that AI training is fair use because their AI models produce original work that is not derived from the sources they use

for training. This is not necessarily true: My reporting has shown that chatbots and image generators can produce near-exact copies of media they were trained on, spitting out near-complete copies of *Harry Potter and the Sorcerer's Stone*, for example, or rendering images that are fuzzy copies of existing artwork. But companies have tried to downplay this fact and focus the copyright discussion elsewhere, even invoking geopolitics and the idea of an international "AI race" as a sort of trump card. "Without fair use access, the race for AI is effectively over. America loses," OpenAI wrote to the Office of Science and Technology Policy last year.

Not everyone in the AI industry is on the same page. Ed Newton-Rex, a former VP of audio at Stability AI, quit his job in November 2023 and wrote on X that, regardless of fair use, which "wasn't designed with generative AI in mind," he didn't see how current AI-training practices "can be acceptable in a society that has set up the economics of the creative arts such that creators rely on copyright." Newton-Rex started a nonprofit called Fairly Trained, which certifies AI models that are trained on properly acquired data.

It's worth noting that Silicon Valley has itself regularly been a victim of IP theft, in the form of software piracy. Partially in response to that problem, major companies have changed how software is distributed. Today, you cannot just buy Adobe Photoshop: Instead, you pay a rental fee to access the program, which verifies your license every time you use it. Microsoft has taken a similar approach with the 365 version of its Office suite, and Google's office software can't be downloaded at all. These companies have made their IP harder to steal by developing new methods of controlling access—an option that is not realistically available to the artists, authors, and open-source-software developers they take material from.

Given the double standard, it's difficult to tell whether Silicon Valley's arguments about fair use are genuine or just legally expedient. On one hand, generative AI is a new technology that raises new questions about the use of copyrighted work. On the other hand, the AI industry's aggressive approach is business as usual for Silicon Valley: moving fast and breaking things. And betting that the lawyers can "clean the mess up."