

Anthropic and Meta Decisions on Fair Use

June 26, 2025

Whether copyrighted works can be freely used to train generative artificial intelligence (“AI”) models is at the core of dozens of lawsuits filed since AI burst onto the scene several years ago. This week, the Northern District of California issued two of the first opinions that begin to answer that question, but there remains a long road ahead before the question is truly settled.

In *Bartz v. Anthropic* (“*Anthropic*”), AI developer Anthropic’s use of copyrighted works to train its large language model (“LLM”) was held to be fair use—but Anthropic’s storage of pirated works was not a fair use, and disputes of fact required a trial on whether Anthropic could maintain a digital library of books it physically purchased. In *Kadrey v. Meta Platforms* (“*Meta*”), the court similarly concluded that training an AI was fair use of copyrighted (and even pirated) works—but strongly suggested the outcome may have been different on a better-developed record.

These courts’ rulings are the first to address the issue of fair use in the context of generative AI specifically and notably contrast with [another recent](#) ruling holding that using a competitor’s data would not be fair use. Especially given the extensive discussion in *Meta* about the limitations on the court’s opinions, these rulings will be far from the last word on the issue of fair use.

Background on the *Anthropic* and *Meta* Litigations

In both *Anthropic* and *Meta*, authors brought suits alleging that LLM developers pirated their copyrighted works and subsequently used them to train the companies’ proprietary LLMs.

Developing LLMs requires developers to obtain pre-existing text, render the text into mathematical representations and then train the LLMs to recognize patterns that enable it to generate responses to user prompts. For example, a user might ask an LLM to write a short story in the style of T.S. Eliot. To fulfill that request, the LLM would

access and process stored representations of Eliot’s writings in its training database to synthesize a new work in the style of Eliot responsive to the user’s prompt.

Anthropic trained its LLM, Claude, using two separate datasets. The first was made up of over seven million pirated works obtained from file-sharing websites, including works authored by the plaintiffs in the instant suit. Anthropic eventually stopped using this dataset and used a second training set by purchasing physical books, scanning them into a digital database, and then destroying the physical copies “for legal reasons.” The *Anthropic* plaintiffs alleged that maintaining both datasets, and using them to train an AI model, were independent acts of copyright infringement.

Meta trained its LLM, Llama, using datasets downloaded from online repositories. One source Meta included in training Llama was pirated “shadow libraries” that provided media—including copyrighted books—for free download without authorization from rightsholders.

Plaintiffs in *Meta* and *Anthropic* filed their initial complaints on July 7, 2023 and August 19, 2024, respectively, in the Northern District of California. Following a contentious discovery period, defendants in *Anthropic* moved for summary judgment and asserted a fair use defense on March 27, 2025. In *Meta*, following a similarly contentious discovery, plaintiffs moved for partial summary judgment on March 10, 2025—arguing in part that defendants’ use of copyrighted works was not fair use. Defendants in *Meta* opposed plaintiffs’ motion and asserted an affirmative fair use defense on March 24, 2025 and filed a cross-motion for summary judgment on March 28, 2025.

Judge Alsup handed down a decision on fair use in *Anthropic* on June 24, 2025, holding that the use of copyrighted works in the training of AI models and the storage of digital copies of scanned purchased books is fair use, but Anthropic’s storage of pirated works involved too many disputed issues of fact to allow the issue of fair use to be resolved at summary judgment.

Just one day later, on June 25, 2025, Judge Chhabria handed down a decision in *Meta*, which cited *Anthropic*. Judge Chhabria held that Meta’s use of copyrighted works to train Llama was a fair use, even where Meta had obtained those works from piracy websites. Notably, however, the *Meta* opinion took pains to explain that its holding was solely based on the record before the court and not a broad holding that all of Meta’s actions were fair use as to all possible plaintiffs.

Using Copyrighted Works as Training Data Was Fair Use...

Both the *Anthropic* and *Meta* opinions ultimately concluded that the AI developers made fair use of authors' works in training their LLMs. But the two opinions took somewhat different paths to reach the same conclusion.

The first fair use factor, **the purpose and character of the use**, favored the AI developers, with both courts agreeing that training an AI model was highly transformative.

The *Anthropic* court analogized training data to human learning and memory, writing that "to make anyone pay specifically for the use of a book each time they read it, each time they recall it from memory, each time they later draw upon it when writing new things in new ways would be unthinkable." The court notably distinguished this case from [Thomson Reuters](#), another AI case in which a Delaware court found that use of copyrighted material as AI training data was not fair use. In *Thomson Reuters*, the plaintiff and defendant created legal research tools with the same purpose and were in direct competition. Anthropic's Claude, on the other hand, had a distinct "purpose and character" from Plaintiffs' books.

The *Meta* court agreed that the use of artistic and literary works to create an LLM was transformative and pointed to Plaintiffs' own testimony that LLMs can be used for a variety of purposes "distinct from creating or reading expressive works" – including obtaining live tax and medical advice or translating documents.

Both the *Meta* and *Anthropic* courts held that the second fair use factor, **the nature of the copyrighted works**, favored Plaintiffs, as all of the copied works, fiction and nonfiction alike, contained expressive elements. And both the *Meta* and *Anthropic* courts held that the third fair use factor, **the amount of the work taken**, favored the AI developers, because copying the books in their entirety was reasonably necessary to successfully train AI models.

The fourth fair use factor, **the effect of the use on the market**, favored the AI developers in both cases, though on exceedingly narrow grounds.

In *Anthropic*, the court reasoned that training data is not public and does not compete in the same market as authors' original works. This led the court to conclude that the "[a]uthors' complaint is no different than it would be if they complained that training schoolchildren to write well would result in an explosion of competing works."

This conclusion notably parted ways with the recent preliminary guidance from the U.S. Copyright Office on the issue of fair use, which argued that “indirect” competition with rightsholders, as well as lost potential licensing revenue, was a harm to the market for their works. The *Anthropic* opinion dismissed both potential harms as being outside of the scope of what the Copyright Act was intended to protect. The *Meta* court agreed that Plaintiffs were not entitled to a market for licensing revenues based on their copyrights but observed that the issue of indirect competition could be a closer call but for the fact that the Plaintiffs had not developed evidence on that point in this case.

The *Anthropic* court did note that its conclusions may be different if the case instead concerned the outputs, which would potentially compete in the same market as the original works. The *Meta* court agreed that the fourth fair use factor could cut against LLM developers if outputs were proven to usurp the market for original works.

The *Meta* court went one step further to emphasize the importance of considering market effects, acknowledging that courts have treated the fourth factor as the “most important factor.” The *Meta* court explicitly rejected the *Anthropic* court’s reasoning on the grounds that if a model could generate works that were similar enough to the original works, they could compete with the original works and indirectly substitute them—thereby diluting the market for the original works. For “when it comes to market effects, using books to teach children to write is not remotely like using books to create a product that a single individual could employ to generate countless competing works with a miniscule fraction of the time and creativity it would otherwise take.”

The *Meta* court went on to note that while the concept of market dilution or indirect substitution is not determinative in many copyright cases, the context of generative AI is different because “no other use—whether it’s the creation of a single secondary work or the creation of other digital tools—has anything near the potential to flood the market with competing works the way that LLM training does. And so the concept of market dilution becomes highly relevant.”

The *Meta* court emphasized the importance of heavily weighing market effects in the context of LLM developers’ use of pirated works. Considering the first and fourth fair use factors, the *Meta* court did not make a similar distinction between pirated and non-pirated copies. Judge Chhabria reasoned that the use of pirated copies, or “shadow libraries,” was relevant to the issue of bad faith, which courts have factored into analysis of the first fair use factor. However, the *Meta* court held that “good faith versus bad faith shouldn’t be especially relevant in the context of fair use” because “whether a given use was made in good or bad faith wouldn’t seem to affect the likelihood of that use substituting for the original.” In other words, if a given allegedly infringing work is unlikely to act as a market substitute of the original work, it is simultaneously more

likely to be sufficiently transformative under the first fair use factor. In the *Meta* court's view, whether Defendants obtained original works through illegal means was not relevant to how likely the infringing work is to usurp or dilute the market of the original works.

The *Meta* court ultimately found in favor of Defendants that training Meta's LLM was a fair use of Plaintiffs' copyrighted works but included extensive dicta explaining the narrowness of its holding. The opinion explained that its conclusion was necessitated by the Plaintiffs' failure to present any empirical evidence that Meta's LLM outputs would harm plaintiffs' ability to profit from their own works—whether through creating substantially similar market substitutes or by promoting the shadow market of pirated works.

... But Maintaining a Data Set Is Another Story

The *Meta* court held that downloading data and LLM training had to be considered as part of the same inquiry and could not be treated as “wholly separate” because such downloading “must still be considered in light of its ultimate, highly transformative purpose: training Llama.” Since Meta's ultimate use of Plaintiffs' books—to train an LLM—was transformative, so too was Meta's downloading of the books. (The court did note that Meta's other actions with respect to the pirated books, like “seeding” them for others to download, were beyond the scope of the opinion and remain an open question.)

The *Anthropic* ruling took a different approach, separating the act of training an LLM from the act of *storing* the training data. Though training an LLM was held to be fair use across the board, the creation and storage of that training was a more fact-dependent inquiry.

Scanning Purchased Books: Fair Use

The *Anthropic* court found that Anthropic's scans of lawfully purchased physical books—which the court characterized as a “mere format change”—was protected by fair use. “Anthropic purchased its print copies fair and square,” the court explained. “With each purchase came entitlement for Anthropic to dispose [of] each copy as it saw fit.” The first fair-use factor favored Anthropic because the format-change from print to digital eased storage and enabled searchability and was thus transformative. The second factor (nature of the works) favored Plaintiffs, but the court accorded it little weight. The third factor (amount of the work taken) favored Anthropic, because copying the works in their entirety was necessary to fulfill the transformative purpose of converting

copyrighted materials from print to digital. The fourth factor was neutral, because a mere change in format did not infringe on any of the copyright holders' exclusive rights.

Notably, however, the *Anthropic* court's conclusion reached only the acts of scanning the books and using them as training data. The court held the record was "too poorly developed" to reach any holding at summary judgment on Anthropic's storage of these books as part of a central library that was retained even after the materials were used as training data.

Storing Pirated Books: Not Fair Use

The *Anthropic* court took a different approach when assessing Anthropic's dataset of pirated works, concluding it was not fair use as a matter of law.

Unlike with the digitized copies, the first fair use factor (purpose and character of the use) favored Plaintiffs, as the copying and storage of digital pirated copies was not transformative. The fact that the stored digital copies were subsequently used to train Claude did not overcome the fact that the copies were first and foremost used to create a library of pirated works. Furthermore, the pirated works were stored indefinitely—both before and after Anthropic used them as training data, even though Anthropic had created a separate dataset with only the scanned copies of the physical books they purchased.

The *Anthropic* court found that the second factor (nature of the works) again favored Plaintiffs due to the expressiveness of their works.

The *Anthropic* court found that the third factor (amount of the work taken) switched to favoring Plaintiffs because unlike the scanned dataset, the pirated dataset was made up of content that Anthropic lacked any entitlement to hold copies of in the first place.

The *Anthropic* court also found that fourth factor favored Plaintiffs. By downloading pirated copies, Anthropic directly reduced the value of Plaintiffs' works though "substitution in their traditional market." The court noted that it was required to "contemplate the likely result were the conduct to be condoned as a fair use—namely to steal a work you could otherwise buy (a book, millions of books) so long as you at least loosely intend to make further copies for a purportedly transformative use (writing a book review with excerpts, training LLMs, etc.), without any accountability. As Anthropic itself suggested, "That would destroy the [entire] publishing market if that were the case."

The Fair Use Defense Going Forward

The *Anthropic* and *Meta* courts' contrasting opinions make clear the fact-intensive nature of the fair use inquiry and the interrelatedness of the fair use factors. The *Anthropic* ruling was notable in its adoption of a bifurcated approach, distinguishing between the use of copyrighted works to train large language models and the methods by which those training datasets are initially compiled. The *Meta* opinion explicitly rejects this bifurcated approach—finding that the compilation of authors' original works is part of the larger process of creating LLMs.

The *Anthropic* court also repeatedly drew on analogies to human behavior—reading, learning, remembering—that other courts may not find as directly parallel to the process of training an LLM. Indeed, the *Meta* opinion pointedly called this an “inapt analogy” in noting that the outcomes were hardly comparable—an educated individual versus a computer program capable of generating “countless” works nearly instantly. These debates over the nature of training and using AI will surely continue, especially with cases challenging the ever more complex ways in which AI generates content like images, audio and video.

The *Meta* opinion offers plaintiffs a roadmap for developing evidence to overcome the fourth fair use factor that we expect other plaintiffs will study closely to shape their discovery requests in other cases. Judge Chhabria explained in detail the types of evidence he expected plaintiffs could develop to show the impact to the market for their works caused by generative AI.

Open Questions

The *Anthropic* and *Meta* decisions only answer one half of the larger legal questions surrounding fair use in the context of generative AI because they only concern AI *inputs* (training data). The question of when or whether AI *outputs* are infringing or fair use remain unanswered, as they were not raised in either case. As *Anthropic* court explained, the “Authors do not allege that any LLM output provided to users infringed upon Authors' works.... Instead, Authors challenge only the inputs, not the outputs, of these LLMs. They point to the fully trained LLMs and the Claude service only to shed light on how training itself uses copies of their works.”

The *Meta* decision provided a partial answer, in dicta, noting that outputs that are substantially similar to Plaintiffs' original works could dilute the market for those original works and harm Plaintiffs' profits—though those facts were not before the court. It accordingly remains to be seen how courts will continue analyze fair use in the

context of generative AI inputs and outputs, which have been asserted in other pending cases, such as *Disney v. Midjourney*.

The *Anthropic* and *Meta* opinions do, however, clearly suggest that the facts surrounding outputs may be relevant to claims based on inputs. Both courts observed that the ability of generative AI to produce exact copies of the plaintiffs' works—which was not in the summary judgment record of either case—could have changed the fair use analysis for AI developers' use of the works as training data. Like the *Meta* opinion's dicta around evidence that may be relevant on the fourth fair use factor, these discussions will likely also serve as a roadmap for other plaintiffs seeking to win a different outcome on fair use in other cases.

These opinions, while significant, are not the end of the road for either case. In *Anthropic*, the parties will proceed to a trial on the pirated works and Anthropic's creation of a digital library. In *Meta*, there is a still-pending motion for summary judgment on the Digital Millennium Copyright Act (DMCA) claims. We also expect to see additional opinions in other AI litigations later this year.

We will continue to follow how these legal and factual questions are resolved as litigation progresses. To stay up to date, subscribe to the Debevoise Data Blog [here](#).

* * *

Please do not hesitate to contact us with any questions.



Megan K. Bannigan
Partner, New York
+ 1 212 909 6127
mkbannigan@debevoise.com



Christopher S. Ford
Counsel, New York
+ 1 212 909 6881
csford@debevoise.com



Abigail Liles
Associate, New York
+ 1 212 909 6387
aeliles@debevoise.com



Joshua Plastrik
Associate, New York
+ 1 212 909 6953
jhplastrik@debevoise.com



Michelle Shen
Associate, New York
+ 1 212 909 6600
mcshen@debevoise.com

