

[IP Litigator, U.S. District Court Issues First Decisions on AI Model Development and Copyright Fair Use, \(Sept. 1, 2025\)](#)

IP Litigator

[Click to open document in a browser](#)

U.S. District Court Issues First Decisions on AI Model Development and Copyright Fair Use

Ryan N. Phelan

Ryan N. Phelan is a registered U.S. patent attorney who counsels and works with clients in all areas of intellectual property (IP), with a focus on patents. As a former technology consultant with Accenture and with a background in computer science and engineering, Ryan has extensive experience in computer system, hardware, and software design, engineering, development and related technologies. He represents numerous startup and Fortune 500 clients with patent matters in technical areas and industries including artificial intelligence (AI) and machine learning, augmented and virtual reality, imaging, medical devices, biometrics data and services, internet and e-commerce, computer networking, data storage and management, encryption and security, mobile telecommunications, consumer electronics, insurance and finance applications, mechanical devices, among others. Ryan is also an adjunct professor at the Northwestern Pritzker School of Law, teaching coursework on patenting software inventions. He is also a frequent author and contributor to Intellectual Property Owner (IPO) and AIPLA white papers and articles, and frequently writes on the subject of AI at his firm's blog PatentNext at patentnext.com.

Recent rulings by two judges in the U.S. District Court for the Northern District of California offer the first merits-based guidance on how “fair use” applies to large artificial intelligence (AI) training, and in particular, language model (LLM) training. These decisions are *Bartz v. Anthropic*, 2025 WL 1741691 (N.D. Cal. June 23, 2025) (referred to herein as “*Anthropic*”) and *Kadrey vs. Meta Platforms*, 2025 WL 1752484 (N.D. Cal. June 25, 2025) (referred to herein as “*Meta*”).

The courts found that using lawfully obtained copyrighted texts for training LLMs can be considered “highly transformative” and can fall under the copyright defense of “fair use,” but that using pirated materials could lead to liability, particularly if the use affects the market for the original works. These rulings shift the legal focus toward the source of training data and whether the AI model’s output causes market harm, setting the stage for future litigation around this issue.

This article provides case overviews of the *Anthropic* and *Meta* cases, explores the four factors of the fair use copyright defense in view of LLM training for each case, and concludes with related implications and takeaways for AI model developers, copyright owners, and AI model end users.

Case Overviews

Bartz v. Anthropic PBC

In *Bartz v. Anthropic PBC*, the court addressed the complex intersection between copyright law and artificial intelligence training. The plaintiffs—authors Andrea Bartz, Charles Graeber, and Kirk Wallace Johnson, along with their affiliated companies—brought suit against Anthropic PBC, an AI firm behind the Claude language model, alleging that Anthropic had unlawfully copied their copyrighted books. Anthropic assembled a massive digital library by both purchasing and pirating millions of books, which it then used to train large language models (LLMs), including Claude.

At issue was whether Anthropic’s various uses of the copyrighted works—including training LLMs, digitizing print copies, using digital pirated copies, and maintaining a central research “library” (a digital database of the

copyrighted books)—qualified as “fair use” under 17 U.S.C. § 107. The court evaluated each use against the four statutory fair use factors and found that while some uses were transformative and thus lawful, others—particularly the use of pirated copies to build a permanent library—were not protected under the fair use doctrine.

Kadrey v. Meta Platforms Inc.

In *Kadrey v. Meta Platforms Inc.*, thirteen prominent authors, including Sarah Silverman and Junot Díaz, filed suit against Meta for allegedly using their copyrighted works—downloaded from unauthorized “shadow libraries”—to train Meta’s large language models (LLMs), particularly the Llama series.

The plaintiffs argued that Meta’s conduct could not qualify as fair use, focusing on harms to the market for their works and the unauthorized nature of Meta’s data acquisition. In contrast, Meta contended that its actions constituted fair use as a matter of law, emphasizing the transformative purpose of LLM training. The court granted summary judgment in favor of Meta, noting the plaintiffs’ failure to adequately substantiate the core theory that Meta’s use would cause significant market harm. However, the ruling applies narrowly to these plaintiffs and does not resolve broader questions about the legality of using copyrighted works in AI training.

Copyright “Fair Use” (Four Factor Analysis by the Courts)

Both the *Anthropic* court and the *Meta* court considered the “fair use” of the copyrighted works. Fair use is a defense typically raised in U.S. copyright disputes and includes an analysis of a four-factor test. Fair use constitutes a defense to allegations of copyright infringement:

[T]he fair use of a copyrighted work ... for purposes such as criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research, is not an infringement of copyright. In determining whether the use made of a work in any particular case is a fair use, the factors to be considered shall include[:]

1. The purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;
2. The nature of the copyrighted work;
3. The amount and substantiality of the portion used in relation to the copyrighted work as a whole; and
4. The effect of the use upon the potential market for or value of the copyrighted work. *Anthropic* at *6.

The following sections consider each of these four factors for both the *Anthropic* and *Meta* cases. In addition, the following sections focus on at least two stages of the AI model development and training process where AI model developers typically face copyright infringement: the first, when the AI model developer stores the copyrighted works in a computer memory for the purpose of training. The second stage is when the trained AI model produces an output – is such output the same or substantially similar to the original copyrighted work or a derivative thereof? For example, for the second of these, a court could focus on whether the output of a given AI model was significantly transformative as opposed to a copy or a derivative work of the original copyrighted material. An AI model can be probed via prompted engineering to determine whether it will output substantially similar works or derivative works from the original copyrighted material. See *Getty Images v. Stability AI*, Case 1:23-cv-00135, (D. Del. Mar. 29, 2003) (Amended Complaint) (Dkt. 13).

Regarding the first of these stages, and as discussed further below, both the *Anthropic* and *Meta* courts were clear that training an AI model with copyrighted works was sufficiently transformative so as to support a fair use defense. In fact, at least according to these two cases, this is one of, if not the most important, factors for finding fair use in each of these cases.

Regarding the second of these stages, in both *Anthropic* and *Meta* cases, the plaintiff-authors failed to allege that output of the respective LLM models produced a same or substantially similar output, and the courts were verbose in highlighting this failure of the authors. That is, had the authors provided additional evidence and arguments regarding a same or substantially similar output from the accused models, then the respective courts indicated that they would have readily (and eagerly) addressed this issue. The authors having failed to raise this

issue, each of the *Anthropic* and *Meta* courts did not rule on the issue and instead highlighted the failure of the authors to do so. We can expect future plaintiffs to address the second of these in a more fulsome manner.

1. The Purpose and Character of the Use

This factor examines whether the use was transformative and whether it served a commercial or nonprofit purpose.

Bartz v. Anthropic PBC

Regarding training LLMs, the court concluded that Anthropic's use of the plaintiffs' works to train LLMs was "spectacularly transformative." Training involved complex processes like tokenization and statistical modeling to teach the LLM to generate new, human-like text. Importantly, the plaintiffs did not allege that the trained Claude system reproduced their works or outputs in a same or substantially similar manner (the hallmark of a copyright infringement claim). The court likened this to a person reading and learning from a book to become a better writer—a transformative use that did not usurp the market for the original works.

Regarding purchased Print-to-Digital book conversion, Anthropic also purchased millions of print books, scanned them, and stored digital copies in its central library. Because each scanned copy replaced its purchased print counterpart, and the digital format merely facilitated internal storage and searchability, the court deemed such use (i.e., a format change of the original purchased works) as contributing favorably to the first factor demonstrating fair use.

In stark contrast, regarding pirated digital book copies, the court found that Anthropic's use of pirated copies to build a permanent, general-purpose library was not transformative. These copies were acquired to avoid "legal/practice/business slog" and were kept indefinitely, even when not used for training. The court emphasized that fair use does not grant AI developers blanket permission to steal and store works simply because some might later be used in transformative ways.

Kadrey v. Meta Platforms Inc.

The first factor—whether the use is transformative and/or commercial—strongly favored Meta. The court found that Meta's use of copyrighted books to train its LLMs served a transformative purpose distinct from the original works. While the plaintiffs' books were intended for consumption as literary or educational texts, Meta used them to extract linguistic patterns and structures to power a tool capable of responding to diverse user prompts.

Even though Meta's ultimate goal was commercial, potentially generating up to \$1.4 trillion in revenue over a decade, the transformative nature of its use was decisive. The court noted that copyright law generally gives more leeway to commercial uses when the new work adds something significantly new. The court also rejected arguments equating LLM training with simple repackaging or copying, noting that Meta's models do not meaningfully output the plaintiffs' original texts. In particular, Meta's LLM was found incapable of reproducing any significant portion of the plaintiffs' copyrighted books, even under conditions designed to provoke memorization. For example, the court noted that Meta's expert employed an "adversarial prompting" technique specifically intended to elicit material from training data, yet no model produced more than 50 tokens (words and punctuation) from the plaintiffs' works. The plaintiffs' own expert achieved similar results in only 60% of tests using the most responsive Llama variant, and further confirmed that Llama was unable to reproduce any substantial portion of the books. Such findings supported the conclusion that Llama could not be used to read or meaningfully access the plaintiffs' copyrighted works.

Further, Meta's controversial use of shadow libraries, while potentially relevant to bad faith, did not outweigh the fundamentally different and transformative nature of the use.

2. The Nature of the Copyrighted Work

This factor considers the creativity and factual nature of the original works.

Bartz v. Anthropic PBC

All of the plaintiffs' books—both fiction and nonfiction—were published and expressive. The court acknowledged that expressive, creative works are closer to the “core” of copyright protection. Because Anthropic specifically valued these works for their expressive qualities in both training and building its library, the court found this factor weighed against fair use across all types of uses—even for those ultimately deemed lawful under other factors.

Kadrey v. Meta Platforms Inc.

This factor favored the plaintiffs. Their works—novels, memoirs, and plays—are highly creative and fall within the heartland of copyright protection. However, courts have historically afforded this factor limited weight, especially when the works have already been published. The court noted that while Meta may not have used the books for their creative expression directly, the statistical patterns it sought to extract were themselves a product of expressive choices like word order, syntax, and style—all protectable elements.

Nonetheless, the court did not view this factor as significantly altering the outcome of the fair use analysis, particularly in light of the highly transformative use under Factor One.

3. The Amount and Substantiality of the Portion Used

Here, the courts assessed whether the amount copied was reasonable in relation to the use.

Bartz v. Anthropic PBC

Regarding training LLMs, although Anthropic copied the entirety of plaintiffs' works for training, the court found this was reasonable given the monumental volume of text required for training effective LLMs. The absence of any public-facing reproduction of plaintiffs' works further supported the finding of fair use.

Regarding purchased Print-to-Digital book conversion, because the digital versions replaced the destroyed print copies and were not shared externally, the court held that copying the entire work was reasonable and aligned with the intended internal use.

In contrast, regarding pirated digital book copies, the court found that copying entire works from pirate sites—particularly to build a centralized research library of indefinite use—was not reasonable. The purpose extended beyond any specific transformative use, and the court noted that almost any level of unauthorized copying would be excessive under these circumstances.

Kadrey v. Meta Platforms Inc.

Although Meta copied the plaintiffs' books in their entirety, the court held that this factor favored Meta due to the necessity of full-text ingestion for the transformative purpose of LLM training. The extent of the copying was deemed reasonable given the technical requirements of training such models. The court emphasized that the key consideration was not the sheer amount of copying, but whether the amount used was excessive in light of the use's purpose.

Given that LLMs perform better with more high-quality data and that partial books would not serve the training purpose effectively, copying entire works was justified and did not weigh against fair use.

4. The Effect of the Use Upon the Market

This factor evaluates whether the use harms the market for or value of the original work. This factor is typically the most critical in a fair use analysis and poses the greatest challenge for the plaintiffs.

Bartz v. Anthropic PBC

Regarding training LLMs, because there was no allegation that Claude's outputs were infringing or substituted for the plaintiffs' books, the court found no adverse market effect. Even potential market competition from LLM-generated works was deemed irrelevant under copyright law, which does not protect authors from generic competition.

Regarding purchased Print-to-Digital book conversion, although Anthropic might have foregone purchasing digital copies, the court found no evidence of redistribution or market usurpation. The internal use of a legally purchased print copy—albeit in a different format—did not harm the existing market in a way actionable under copyright law.

Regarding pirated digital book copies, this use had a direct and deleterious effect on the market. By copying works it could have lawfully purchased, Anthropic displaced market demand on a copy-for-copy basis. The court emphasized that permitting such behavior would effectively destroy the publishing industry, as it would incentivize theft in the name of downstream transformative use.

Kadrey v. Meta Platforms Inc.

The court identified three potential types of market harm: (1) regurgitation of the original works, (2) loss of licensing revenue for AI training, and (3) market dilution through proliferation of similar AI-generated content.

The first two arguments failed due to insufficient evidence. Llama was not capable of meaningfully regurgitating the plaintiffs' works, and courts do not recognize a right to licensing revenue for transformative uses. While the third argument—market dilution—was conceptually strong and could be highly relevant in future cases, the plaintiffs failed to plead or support it with evidence. Thus, they could not create a triable issue of fact on this point.

The court stressed that while market dilution from AI-generated content may be a valid concern under copyright law, it must be substantiated with evidence. As such, Factor Four also favored Meta.

Court's Conclusions and Takeaways

Bartz v. Anthropic PBC

The *Anthropic* court's overall analysis reflected a nuanced application of the fair use doctrine. It recognized fair use for the training of LLMs using copyrighted books, which was considered transformative. So was the scanning of purchased print copies for internal digital storage and use.

However, the *Anthropic* court denied the fair use defense for the use of pirated copies for building a central research library, which was not considered transformative and failed all four fair use factors.

Nonetheless, the *Anthropic* court granted summary judgment in favor of Anthropic for the training and format conversion uses, but denied it for the pirated library copies. The case is set to proceed to trial to determine liability and damages for the unauthorized acquisition and retention of those pirated materials.

This decision reinforces that while AI development may qualify for fair use under certain conditions, courts will scrutinize the methods and intentions behind data acquisition—especially where piracy is involved. AI innovators must balance transformative use with lawful sourcing to stay within the bounds of copyright law.

Kadrey v. Meta Platforms Inc.

The ruling in *Kadrey v. Meta Platforms Inc.* offers a nuanced but limited precedent. While Meta prevailed on summary judgment, the court's decision hinged on the plaintiffs' failure to develop and present a compelling case on the most critical issue—market harm. The decision does not validate Meta's use of copyrighted works in AI

training as lawful per se; rather, it underscores the importance of presenting the right evidence under the fair use framework.

This case may serve as a roadmap for future litigants—highlighting the potential viability of market dilution arguments and signaling that courts remain receptive to fair use challenges in the context of transformative AI technologies, so long as they are properly developed and supported.

Also, being the second of the two cases, the *Meta* court voiced its differences and concerns with the *Anthropic* court, stating that the *Anthropic* court “focused heavily on the transformative nature of generative AI while brushing aside concerns about the harm it can inflict on the market for the works it gets trained on.” *Id.* at *11. The *Meta* court took issue with the *Anthropic* court’s reasoning that “[s]uch harm would be no different ... than the harm caused by using the works for ‘training schoolchildren to write well,’ which could ‘result in an explosion of competing works.’” Instead, the *Anthropic* court was sympathetic to the plaintiff-author’s concern regarding market harm: “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. This inapt analogy is not a basis for blowing off the most important factor in the fair use analysis.” *Id.*

Conclusion

The court decisions involving *Meta* and *Anthropic* mark the beginning of what is expected to be a wave of legal rulings addressing copyright issues in generative AI. While these initial cases centered on large language models (LLMs) trained on books, future outcomes may vary depending on the nature of the training data and output. Notably, cases involving image-based content like *Getty v. Stability AI* or code-based output, such as in *Doe 1 v. GitHub, Inc.*, No. 4:22-cv-06823-JST (N.D. Cal.), may yield different legal analyses, highlighting the evolving complexity of copyright law as applied to various AI-generated modalities.

For example, such cases also explore the important question: what type of relationship should copyright holders have with AI model developers? Such a question is not only important to authors of books, articles, and other such written materials, but is also important to implications regarding computer software code, which is the foundation of most companies’ IP. For example, if an AI tool is used to create valuable source code for a company’s product or service, who owns that source code (if anyone per the authorship requirements under U.S. copyright law), and is that source code subject to potential copyright violations for using or generating the same or substantially similar code from which the AI tool was trained from?

Implications for Artificial Intelligence (AI) Model Developers. For AI model developers and related stakeholders—especially tech platforms, cloud providers, publishers, and data brokers—these decisions can signal a need for immediate action. Organizations should consider auditing their training datasets and vendor agreements to ensure all source materials are lawfully obtained, carefully document any market impact, and update internal policies accordingly. Legal and technical leaders should consider collaborating closely to align data practices with emerging legal expectations. For example, one approach from the *Meta* and *Anthropic* decisions can involve digitizing legally purchased physical books and then destroying the originals. To further reduce legal exposure, LLM developers can implement output guardrails that prevent or minimize the reproduction of copyrighted content.

Implications for Copyright Owners. Copyright owners may want to keep sensitive data a trade secret. If desired, a copyright owner seeking to license its private data for training purposes may want to consider doing so under a license agreement that includes privacy restrictions pursuant to a non-disclosure agreement (NDA) to prevent the data from leaking to the public. One of the main issues for copyright owners in the *Anthropic* and *Meta* case was that the copyrighted works were public, such that the authors could not control their use for training. This will always be the case for books and other copyrighted works intended for public consumption. But for trade secret data, such as proprietary datasets, more control can be exercised to monetize valuable datasets for AI training.

Implications for AI Model Users. Companies utilizing large language models (LLMs) can take key measures when contracting with LLM developers. First, they should consider auditing the training data by requesting a comprehensive list of datasets used to train or fine-tune the model, ensuring no pirated content from shadow libraries is included. Second, they could also consider verifying that the LLM incorporates effective guardrails to prevent the output of copyrighted material, with internal testing by creative staff to confirm their effectiveness. Finally, companies should consider negotiating strong indemnification provisions to protect against potential copyright infringement claims, recognizing that while current litigation has focused on developers, users may still face some legal exposure.

We can expect appeals from these cases and the appellate courts to take up these issues and provide guidelines. However, this could take several years, and these issues will likely find their way to the Supreme Court for ultimate resolution. This assumes, of course, that Congress does not act first to provide a statutory framework.

DISCLAIMER: The information contained in this article is for informational purposes only and is not legal advice or a substitute for obtaining legal advice from an attorney. Views expressed are those of the author and are not to be attributed to Marshall, Gerstein and Borun LLP or any of its former, present, or future clients.