

Research Article

A Study on the Judicial Application of Fair Use for Artificial Intelligence Training Data in the United States

Xiaojing Qin

School of Law, Beijing Normal University, China

***Related declarations are provided in the final section of this article*

ARTICLE INFORMATION

Received: 09 Feb 2026

Accepted: 18 May 2026

Published: 14 Sep 2026

Copyright © 2026 The Author(s). The Authors retain copyright. Licensed under the [Creative Commons Attribution 4.0 International License \(CC BY 4.0\)](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

The rise of generative artificial intelligence has made the fair use determination of training data a cutting-edge controversy in copyright law. This article systematically examines the judicial application experience of U.S. federal courts in the cases of Thomson Reuters v. ROSS Intelligence, Bartz v. Anthropic, and Kadrey v. Meta. The study finds that the divergences among the three cases in the weighting of the four factors and the conclusions of their rulings are not logical contradictions, but rather reveal the flexible application mechanism of the fair use doctrine under different technological scenarios: the functional relationship between the AI product and the original work constitutes the core variable determining the finding of transformative use; the legality of the data source constitutes an independent factual basis for negating fair use; and the market impact analysis presents a typological structure ranging from direct substitution, licensing market harm, to market dilution.

Keywords: Artificial intelligence training data; Fair use; Transformative use; Market dilution; Four-factor analysis framework

INTRODUCTION

The institutional evolution of copyright law demonstrates that each generational leap in information dissemination technology has profoundly challenged the interpretive boundaries of fair use rules, from the Sony Betamax case triggered by video recording technology, to the fair use determination of large-scale digital copying in the Google Books case in the Internet era.[11,9]The rise of generative artificial intelligence is opening the latest chapter in this institutional evolution: when large models "digest" humanity's existing intellectual achievements at the scale of billions of parameters, whether the automated scraping and statistical modeling of massive copyrighted works during training constitutes an infringement of copyright holders' exclusive rights, or may be exempted through the fair use doctrine, has become one of the most cutting-edge controversial issues in the field of intellectual property law.

The core institutional tool provided by U.S. law to address this issue is the four-factor analytical framework established under Section 107 of the Copyright Act: the purpose and character of the use, the nature of the copyrighted work, the amount and substantiality of the portion used, and the effect of the use upon the potential market for or value of the copyrighted work.[17,5] This article systematically examines the judicial application experience of the U.S. four-factor analysis framework in the field of AI copyright.

JUDICIAL PRACTICE OF THE U.S. FAIR USE DOCTRINE

The fair use doctrine established in Section 107 of the U.S. Copyright Act, with its open four-factor analytical framework, distinguishes it from the enumerated limitation model of civil law systems, and has historically played a leading role in the formation of rules in copyright disputes arising from new technologies. Between 2023 and 2025, U.S. federal courts made the first substantive rulings on the fair use issue of AI training data. Due to the differences in technology types and the divergence in rulings among these three groups of cases, they collectively constitute core samples for examining the boundaries of the fair use doctrine's application in the AI era. The Thomson Reuters case (D. Del., February 2025) involved the training use of a non-generative AI legal search engine on Westlaw case summaries, and the court ultimately denied the fair use defense, with the core reason being that the AI product constituted direct competitive substitution with the original work in the same functional market. The Bartz case (N.D. Cal., June 2025) involved the training use of the generative AI large language model Claude on literary works, and the court found the training use to be "extremely transformative," holding that fair use was established, while simultaneously explicitly excluding the application of fair use to training data obtained through pirated channels. The Kadrey case (N.D. Cal., June 2025) similarly involved the training use of the generative AI large language model LLaMA, and while the court found fair use, it creatively proposed a "market dilution" theory, extending the fourth factor analysis to the indirect competitive effect of AI's large-scale generation of works of the same type on the market space of the original works.⁴ The three cases exhibit a complementary distribution in terms of technology types: Thomson Reuters involves non-generative AI, while Bartz and Kadrey involve generative AI; they form a sharp contrast in terms of rulings, with the former denying fair use and the latter two affirming it; and each has a different analytical focus, with Thomson Reuters taking market substitution as the decisive factor, Bartz centering

its argument on transformative use, and Kadrey expanding the theoretical dimension of market impact. This difference is not a contradiction in judicial reasoning, but rather reveals the flexible application mechanism of the fair use four-factor analysis under different technological scenarios. The functional positioning of the AI product, its competitive relationship with the original work, and the method of obtaining training data collectively constitute the key factual variables affecting the weight allocation of each factor. The following analysis proceeds in sequence.

THE THOMSON REUTERS CASE

1. Case Background and Institutional Significance

The judicial application of the U.S. fair use doctrine in the field of AI copyright is centrally manifested in three landmark cases adjudicated by federal courts between 2023 and 2025. This section sequentially examines the Thomson Reuters case, the Bartz case, and the Kadrey case, which respectively involve two technology types—non-generative AI and generative AI—and exhibit significant divergences in the weighting of the four factors and final conclusions, collectively constituting a diverse judicial spectrum of U.S. federal court responses to the fair use issue of AI training. The following case-by-case analysis proceeds in chronological order.

The Thomson Reuters case is one of the first cases in which a U.S. federal court made a substantive ruling on the fair use defense for AI training data. [13,14] The plaintiff, Thomson Reuters, operates the legal research platform Westlaw, which features attorney-edited summaries of judicial opinions (headnotes) organized by a "key number system." [14] The defendant, ROSS, is an AI startup company that developed a legal research tool based on natural language search. After ROSS sought and was denied a content license from Westlaw, it turned to a third-party company, Legalese, to create a training dataset of approximately 25,000 question-and-answer pairs, which the plaintiff alleged substantially copied the case summaries. [14] This case involved non-generative AI, and the court explicitly found that ROSS "is not a generative AI" and that its product did not generate new content but rather retrieved and presented existing judicial opinions. [14] This technological feature directly determined that the functional competitive relationship between ROSS's product and Westlaw was closer than that of generative AI, thereby profoundly affecting the court's analysis of factors one and four.

2. Fair Use Four-Factor Analysis

2.1. Factor One: The Interaction Between Commercial Nature and Transformative Use

The analysis of factor one in this case presented itself as an interactive balancing between commercial nature and transformative use, [13] and ultimately revealed a judicial logic of general significance for subsequent cases: the determination of transformative use is closely related to the functional positioning of the AI product.

With respect to commercial nature, the court found that ROSS's use was "undoubtedly commercial in nature" and that "one of ROSS's goals was to compete with Westlaw." [13,14] However, the court refused to accord decisive weight to commercial nature: citing *Google v. Oracle*, it noted that the Supreme Court, in technology-related cases, "placed greater weight on transformativeness than on commerciality," and that "many undisputedly fair uses are commercial in nature." [13,14,3] This position delineates the normative status of

commercial nature in the fair use analysis of AI: it is a modifying factor affecting the conclusion of factor one, rather than an independent negating factor.

Transformative use was the core of the dispute in this case, and the opposition between the parties essentially reflected a structural difficulty in the fair use determination of AI training: should the same technological process—converting textual data into numerical representations for machine learning—be characterized as the direct exploitation of the original work's creative expression, or as "intermediate copying" for the purpose of creating a new product? The plaintiff argued that ROSS merely "translated" the case summaries "into numerical data," constituting a "typical derivative work," the very purpose of which was "to replicate the creative decisions of Westlaw's attorney-editors." [13] The defendant, citing the "intermediate copying" precedents of *Sega* and *Sony*, argued that encoding text into numerical data was merely an intermediate step in developing a "new product," and that the final product did not contain infringing material. [13,10]

In its 2023 memorandum opinion, the court did not rule directly but established a conditional analytical framework: if the defendant's account were true—that the use was merely an intermediate step in creating a new product and the final product did not output infringing material—then it would constitute "transformative intermediate copying"; if the plaintiff's account were true—that the use was intended to replicate Westlaw's creative editorial work—then the intermediate copying analogy would not apply. [14] However, in its 2025 renewed ruling, the court's position fundamentally shifted: it found that ROSS's use did not constitute transformative use, with the core reason being the direct functional competitive relationship between ROSS and Westlaw and ROSS's failure to prove that its product achieved a substantial functional transformation [14]. From the 2023 "conditional reservation" to the 2025 "explicit negation," this judicial trajectory establishes a rule of general guiding significance: when the AI product constitutes direct competition with the original work in the same functional market, the "intermediate copying" argument at the technical level cannot support a finding of transformative use.

2.2. Factors Two and Three: Limited Weight but Analytical Framework of Referential Value

With respect to the second factor, the court found that while the case summaries involved a certain degree of creative selection, the editorial creativity was less than that of a novelist or artist creating a work from scratch, and the nature was closer to informational arrangement. [13] The court accordingly found that the second factor favored the defendant, while simultaneously noting that "this factor rarely plays a significant role in fair use disputes." [13,9,14] With respect to the third factor, because ROSS's final output product presented to end users was the judicial opinions themselves rather than Westlaw's case summaries, the protected expression displayed to the public was limited. [14,9,13] The 2025 ruling found that both factors two and three favored the defendant. [14] Although both factors favored the defendant, their weight was insufficient to overturn the conclusions of factors one and four.

2.3. Factor Four: Typological Analysis of Market Impact

Factor four was accorded the greatest weight by the court in this case, and its analysis presents a typological structure ranging from market substitution to licensing market impact to indirect competition, which provides the basic paradigm for market impact analysis in subsequent AI copyright cases.

The market substitution dimension constituted the core of the factor four analysis in this case. The court engaged in careful conceptual distinction in its 2023 ruling: although ROSS and Thomson Reuters were competitors in the legal research market, "that alone does not reveal whether ROSS is a substitute for Westlaw," and if the use is transformative, "then it is not a market substitute." [13] The court cited a theoretical proposition that was repeatedly applied in subsequent cases: "the more the copying is for a purpose different from that of the original, the less likely it is that the copy will serve as a satisfactory substitute for the original." [13,9,2] However, in its 2025 renewed ruling, the court changed its position, finding that ROSS intended to compete with Westlaw by developing a market substitute, and explicitly stated that "copyright encourages people to develop things that are beneficial to society... its creators are entitled to be compensated for it." [14] Market substitution thus became the decisive force negating fair use.

With respect to licensing market impact, the plaintiff argued that there existed a traditional licensing market for Westlaw data and an emerging AI training data market, and that ROSS's payment of over one million dollars to Legalese for training data proved the existence and harm of that market. [13] The 2025 ruling adopted this argument, finding that ROSS's conduct had an adverse impact on Thomson Reuters's potential derivative licensing market, namely the AI training data market. [14]

With respect to the indirect competition dimension, the court candidly acknowledged the difficulty of policy judgment in its 2023 ruling and issued a cautious signal regarding the broader policy implications of AI training. [13] This cautious stance reflects the court's institutional self-awareness when facing the emerging field of AI copyright.

3. Comprehensive Conclusion and Adjudicative Implications

The 2025 ruling ultimately denied the fair use defense. In the comprehensive balancing of the four factors, the court found that factors one and four together favored the plaintiff with decisive weight, while factors two and three, although favoring the defendant, were of limited weight, and the combined effect of the former two outweighed the latter two. [14]

The adjudicative significance of this case can be distilled from the following three levels. First, with respect to the relationship between technology type and fair use, the functional competition between a non-generative AI search engine and the original work is more direct than that of generative AI, and the subsequent Bartz case, in addressing large language models, adopted a distinctly different analytical path. Second, with respect to the structure of transformative use determination, this case reveals three core dimensions of judgment: whether the AI product serves a purpose different from the original work, whether it contains the protected expression of the original work, and whether the use of the original work is merely an intermediate step to

achieve a transformative purpose. Third, with respect to market impact analysis, the typological structure of market substitution, licensing market impact, and indirect competition established in this case provides an analytical paradigm of referential value for subsequent cases.

THE BARTZ CASE

1. Case Background and Theoretical Positioning

The Thomson Reuters case, with "functional competitive relationship" as the core variable, denied the fair use defense for non-generative AI. The theoretical question that follows is: when the technological architecture of AI shifts from "retrieving existing content" to "generating new content," how will the fair use analytical framework be adjusted? The Bartz case obtains its institutional coordinates within this theoretical lineage. This case is the first time a U.S. federal court made a substantive ruling on the fair use issue of generative AI training,[21]and its core contribution lies in revealing a proposition of general theoretical significance: the determination of transformative use does not depend on the form of the technological process itself, but on the functional relationship between the purpose of the particular use and the expressive purpose of the original work. When the two are "sufficiently orthogonal," transformative use may be established even if the training process involves full copying of the original work.

In this case, the defendant, Anthropic, is the developer of the large language model Claude. Its acquisition of training data involved three methods: downloading over 7 million pirated books from pirate websites such as Books3 and Library Genesis, with the founder explicitly acknowledging that piracy was chosen to avoid "legal/practical/commercial hassle";[21]purchasing large quantities of physical books and conducting destructive digital scanning;[21] and consolidating the above sources into a central library planned to "permanently store all content." [21]The most critical factual premise of this case lies in the state of the output end: the plaintiffs did not allege that Claude provided any content infringing their works' copyright to users.[21]

Judge Alsup's methodological innovation in the analytical framework endowed this case with theoretical value beyond individual adjudication. The court divided Anthropic's conduct into three distinct "uses"—the copying used to train the LLM, the copying converting legally purchased physical books into digital format, and the copying downloading from pirated sources and retaining in a central library—and applied the fair use four-factor analysis to each separately.[21]The court cited Warhol to confirm that the determination of "use" depends on "an objective examination of what the user actually did with the work" rather than subjective intent.[21,18]This "conduct differentiation" methodology means that different uses of the same work by the same actor, due to their respective purposes and effects, can yield differentiated fair use conclusions, and this refined analytical approach provides an important methodological paradigm for subsequent judicial practice.

2. Fair Use Four-Factor Analysis

2.1. Factor One: Theoretical Deepening of Transformative Use Determination

Factor one is the area of greatest theoretical contribution in this case. The court reached affirmative conclusions for two of the three uses and a negative conclusion for one, thereby advancing the "functional relationship" analytical structure established in *Thomson Reuters* to a more refined theoretical level: whether transformative use is established depends not on some general characteristic of AI technology, but on the functional distance between each specific use and the expressive purpose of the original work.

With respect to the training copying conduct, the court made the most forceful finding of AI training trans formativeness in U.S. judicial practice to date: "the purpose and character of using copyrighted works to train an LLM to generate new text is classic transformative use." [21] This finding was based on the successive rejection of three opposing arguments. In response to the argument that "AI training is equivalent to human learning and authors have the right to exclude," the court began from the historical tradition of the copyright system, noting that for centuries, readers' "reading, appreciating, remembering, and internalizing" the grand themes and stylistic solutions of works had been self-evidently lawful conduct. [21] The theoretical gist of this argument is: copyright law protects specific expressions, not the social circulation of knowledge, methods, and styles carried by works, and AI training functionally belongs to the latter category. In response to the argument that "training is intended to memorize creative elements," the court cited Section 102(b) to confirm that copyright does not extend to "procedures, processes, systems, methods of operation, concepts, principles, or discoveries," and further noted that the process by which an LLM distills grammar, composition, and style from thousands of works is essentially no different from the fusion by human authors of imitating classic writing. [21,16] Of particular note is the court's active distinction from the *Thomson Reuters* case. Judge Alsup pointed out that the *Thomson Reuters* case involved "not generative AI," and its training purpose was to develop a search tool competing with Westlaw in the same functional market, "that was not transformative"; by contrast, the use in this case was "sufficiently [orthogonal] to anything the copyright owner could legitimately expect to control." [21,13,19] The court characterized the essence of generative AI training conduct as "turning a sharp corner and making something different." [21] Thus, the divergence between the two cases on factor one receives precise theoretical localization: *Thomson Reuters* negates trans formativeness due to the functional homogeneity between the AI product and the original work; *Bartz* affirms trans formativeness due to the functional orthogonality between the AI training purpose and the expressive purpose of the original work. The two are not contradictory, but together establish a continuum with "functional relationship" as the core variable—from functional homogeneity (trans formativeness not established) to functional orthogonality (trans formativeness established)—and the positioning of the AI product on this spectrum determines the analytical conclusion of factor one. With respect to format conversion after lawful purchase, the court found that it constituted fair use. The court focused on the key fact that the physical original was destroyed, "one copy replaced the other, and the digital copy was not displayed, shared, or sold to anyone outside." [21] This finding theoretically means that when the copying is not publicly disclosed or displayed, transformative use is established.

With respect to the pirated library copy, the court's analytical stance underwent a fundamental shift. Judge Alsup ruled: "the pirated copy to build a research library without paying for it is itself a use, and it is not transformative." [21] The court explicitly rejected Anthropic's strategy of subsuming the pirated acquisition conduct under the "transformative halo" of the subsequent training use, and further cited Warhol to note that "the same copy can be used in one way and then in another." [21,18] The court further distinguished the Sega precedent, in which the defendant "purchased commercially available copies of game cartridges" before copying, whereas in this case "there was no authorized copy in existence." [21,20] The court accordingly concluded: "downloading from pirate websites works that could have been purchased is inherently, irredeemably infringing." [21] This ruling forms an inter-case resonance with the Thomson Reuters case on the issue of data source legality. Although the two cases fundamentally diverged on transformativeness and market impact, they reached a clear consensus on the proposition that the illegality of data acquisition will negatively affect or even completely preclude the establishment of fair use.

2.2. Factors Two and Three: High Protection for Expressive Works, Full Copying, and Application of the Public Exposure Standard

The court applied the second factor analysis uniformly to the three types of use, finding that this factor militated against fair use for all copies. [21] Unlike the informational arrangement involved in Thomson Reuters, this case involved literary works with rich expressive elements, and "they were selected precisely because of their expressive qualities" for training. [21] The second factor accordingly favored the plaintiff to a significantly greater degree than in Thomson Reuters. However, the court simultaneously noted that the second factor's "primary function is to help evaluate the other factors," [21,2] and does not have independent decisive weight in the overall analysis—a judgment consistent with Thomson Reuters.

With respect to the training copying conduct, the court found that although Anthropic copied the entirety of the works, the third factor nevertheless favored fair use. This seemingly paradoxical conclusion was based on the court's application of the Google Books standard: "what matters is not 'the amount and substantiality of the portion used' in making the copy, but the amount and substantiality of what is thereby made available to the public that could serve as a competitive substitute for the original use." [21,9] Since there was no traceable connection between Claude's output and the plaintiffs' works, the full copying during training was "accordingly especially reasonable." [21] In response to the plaintiffs' argument that "copying was not strictly necessary," the court established an important distinction: reasonably necessary does not mean strictly necessary. [21] The court found that Anthropic's need for billions of words per model and the "compelling benefit" of training on "the strongest examples of writing" were "not offset by revealing any part of the works themselves to the public." With respect to the pirated copy, the court found that the third factor disfavored fair use. [21]

The difference in reasoning paths between the two cases on the third factor is worth noting: Thomson Reuters focused on the compositional proportion of training data in the AI's final product, while Bartz focused on whether the AI product exposed substantial content of the original works to the public. The latter path is theoretically more precise because it anchors the third factor analysis to the core interest protected by copyright law—preventing unauthorized

dissemination of protected expression to the public—rather than a pure statistical accounting of training data.

2.3. Factor Four: Fundamental Divergence in Market Impact Theory

Factor four is the area where this case and the Thomson Reuters case produced the deepest opposition in judicial philosophy, and is also the focal point of theoretical tension between this case and the subsequent Kadrey case.

With respect to market substitution, the court took the factual state of the output end as the decisive basis: "the plaintiffs concede that training the LLM has not resulted in any exact copies or even infringing imitations of their works being provided to the public. If that were not the case, it would be a different case." [21] In Thomson Reuters, ROSS was found to be a market substitute for Westlaw, with the core being the functional overlap between the two in the legal research services market; in Bartz, the plaintiffs failed to prove that Claude output content substantially similar to their works. This factual difference is the fundamental reason for the diametrically opposite conclusions on factor four in the two cases, and it reaffirms the status of "functional relationship" as the core variable in fair use analysis.

With respect to licensing market impact, the court made a ruling of profound theoretical significance. The court assumed that a licensing market "might develop," but ruled that "even if so, that market is not one that copyright law entitles authors to exploit." [21] The court positioned the training use as a fair use area that copyright holders "cannot legitimately expect to control." [21] This stance fundamentally opposes the philosophy of the Thomson Reuters case, and the divergence can be summarized as two distinct legal philosophical orientations: the Thomson Reuters case follows the logic of "where there is a market, there is a right"—the existence or potential existence of a licensing market itself constitutes proof of a protectable interest of the copyright holder; the Bartz case follows the logic of "where there is a right, there is a market"—only uses over which copyright law grants the copyright holder exclusive control can make licensing market harm legally significant. This latter stance profoundly reflects the economic foundation of the fair use doctrine as a market failure remedy: when transaction costs are too high (piecemeal licensing of millions of works is technically and economically nearly infeasible), fair use intervenes as an alternative mechanism for interest allocation, and its justification lies precisely in the failure of the licensing market.

However, the long-term stability of Judge Alsup's position warrants careful assessment. In recent years, major AI model companies have begun to adopt licensing models on a large scale, such as OpenAI's data licensing agreement with Shutterstock in 2023. [12] The increasing prevalence of licensing practices is fundamentally reshaping the factual premise of "whether a training data licensing market exists." If such market practices further develop and mature, whether courts can continue to maintain the judicial position that "copyright holders have no right to control training uses" will become one of the most controversial theoretical issues in the development of U.S. AI copyright law.

With respect to indirect competition, the plaintiffs argued that training LLMs would lead to a proliferation of works competing with their works. The court assumed this argument to be true, but responded with a highly controversial analogy: "the plaintiffs' complaint is no different from complaining that training schoolchildren to write good essays will lead to a proliferation of competing works. That is not the kind of competitive or creative substitution that copyright law is concerned about. The Act is designed to promote the creation of original works, not to protect authors from competition." [21,20] This analogy was directly criticized by Judge Chhabria in the subsequent Kadrey case. [6] The latter pointed out that using books to train children to write is entirely different from using books to create a product that a single individual can use to generate countless competing works with minimal time and creativity. This open disagreement between different judges of the same court profoundly reflects the cutting-edge controversial nature of the market dilution theory in the field of AI copyright, which will be more systematically examined in the analysis of the Kadrey case below.

3. Comprehensive Conclusion and Theoretical Distillation of the Comparison Between the Two Cases

The court ruled separately according to the three types of use: the training copying conduct constituted fair use, and "the technology at issue is one of the most transformative technologies many of us will see in our lifetimes"; [21] the pirated library copy did not constitute fair use, and "every factor points against fair use." [21]

In the allocation of factor weights, for the training copying conduct, factors one, three, and four all favored the defendant, with only factor two favoring the plaintiff but of limited weight; for the pirated copy, all four factors consistently pointed to the plaintiff. Juxtaposing the two cases for examination, their divergence can be distilled at three theoretical levels. First, with respect to the structure of transformative use determination, the two cases together establish a spectrum of judgment with "functional relationship" as the core variable: from functional homogeneity in Thomson Reuters (transformative use not established) to functional orthogonality in Bartz (transformative use established), which determines the direction of factor one analysis. Second, with respect to the philosophical orientation of market impact analysis, the opposition between the two cases of "where there is a market, there is a right" and "where there is a right, there is a market" touches upon the deepest normative foundation of the fair use doctrine: should the boundaries of copyright exclusivity be pre-defined by law, or dynamically delineated by market practice? This philosophical divergence will continue to influence the trajectory of U.S. AI copyright adjudication. Third, with respect to data source legality, the two cases reached a clear consensus across all other divergences: the illegality of the method of data acquisition constitutes an independent and sufficient factual basis for negating fair use.

THE KADREY CASE

1. Case Background and Theoretical Positioning

On June 25, 2025, Judge Chhabria of the U.S. District Court for the Northern District of California issued a partial summary judgment ruling in the Kadrey case, denying the plaintiffs' partial summary judgment motion, granting Meta's cross-motion for summary judgment, and

finding that, based on the specific evidentiary record in this case, Meta's use of the plaintiffs' works to train its large language model Llama constituted fair use.[6]

The theoretical contribution of this case's ruling lies not in its fair use conclusion itself—the court itself limited the precedential effect of that conclusion with unusually candid language—but in Judge Chhabria's systematic analysis of factor four, which for the first time at the judicial level affirmatively constructed the analytical category of "market dilution," bringing the indirect market competition triggered by AI training within the scope of fair use review, and openly criticizing the core argument of Judge Alsup in the Bartz case. From a comparative law perspective, the proposition revealed by this case—that "highly transformative use and substantial market harm can coexist"—has direct methodological referential value for Chinese courts in addressing AI training fair use issues under the integrated "four-factor—three-step test" framework in the future. The following analysis proceeds factor by factor in sequence.

2. Fair Use Four-Factor Analysis

2.1. Factor One: Purpose and Character of the Use—Transformative Use Determination and Its Boundaries

First, the determination of transformative use. Judge Chhabria characterized factor one as "clearly favoring Meta," with the core of his reasoning being the fundamental divergence in purpose: the inherent purpose of the plaintiffs' books was "to be read by readers for entertainment or education," while Llama's functional positioning pointed to "fundamentally different" scenarios such as drafting text, translating languages, and generating code.

[6] The court, citing the transformative use logic in Oracle of "creating a new platform that programmers could use conveniently," found that LLM training likewise had an independent functional purpose separate from the original work. [3,6] The court noted that what Meta created was a general-purpose platform usable by any user, and that its copying had the potential to "exponentially expand creative expression." [6]

Second, commercial nature and data acquisition methods. The court candidly acknowledged that Llama's development had significant commercial motivation, and cited Meta's own revenue projections as support [6] However, in the context of highly transformative use, the adverse weight of the commercial factor correspondingly diminished. [7,6] With respect to data acquisition methods, the court rejected, on circular reasoning grounds, the plaintiffs' argument that Meta's downloading from "shadow libraries" such as LibGen should be directly characterized as "piracy" so as to preclude the possibility of fair use. [6] Of note is that the court here explicitly drew a methodological boundary with the Bartz case: Judge Alsup had severed the downloading conduct from the training conduct for independent fair use evaluation; [21] Judge Chhabria argued that the downloading conduct should be evaluated holistically together with its ultimate "highly transformative" training purpose, to avoid artificially severing the various links of an integrated chain of use. [6,9] The essence of this divergence is not a difference in factual findings, but a methodological choice regarding how the "unit of analysis" in fair use analysis should be delineated—whether to evaluate each physical copying act separately, or to measure holistically according to the ultimate purpose of use—the answer to which will directly determine the legal characterization of each link in the entire AI training chain.

2.2. Factors Two and Three: Ancillary Analysis of Subsidiary Factors

With respect to factor two, the court found that this factor pointed against fair use. Judge Chhabria cited *Hachette Book Group, Inc. v. Internet Archive* to confirm that both fictional and non-fictional works lie at the core of copyright protection.[4,6]In response to Meta's argument that "LLMs merely extract functional elements from books," the court explicitly rejected it, noting that unlike in *Sega Enterprises Ltd. v. Accolade, Inc.* and other computer code cases where copying was merely to access non-protected functional elements, the "word order, word choice, grammar, and syntax" that LLM training relies upon precisely carry the author's creative expression; even if what the LLM "learns" is only "statistical relationships," these statistical relationships themselves are precisely the product of creative expression.[6]The significance of this holding is that it blocks the interpretive channel by which AI companies might evade the adverse implications of factor two under the guise of "functional extraction." However, the court simultaneously noted, citing existing precedents, that factor two "rarely plays an important role in fair use disputes," and its impact on the overall balancing is limited. [6,9]

With respect to factor three, although the court confirmed that Meta made full-text copies of the plaintiffs' books, it nevertheless found that this factor favored Meta. The court's reasoning followed two paths: first, in terms of quantity, full-work training could better improve LLM performance than partial training, and complete copying was reasonably necessary for the training purpose;[8,6]second, in terms of substantiality, given that Llama could not output any meaningful amount of text from the plaintiffs' works, in weighing this factor the focus should not be on how much the copier used, but on how much copyrighted material it made available to the public.[6,4]However, the court candidly acknowledged in a footnote that there was an as-yet-unexplored inherent tension here: the more fully trained the LLM, the stronger its ability to generate competitive works, and the correspondingly higher the risk of market dilution—the release of this tension would have to be taken up by the analysis of factor four.

2.3. Factor Four: The Tripartite Structure of Market Harm and the Affirmative Construction of "Market Dilution" Theory

Judge Chhabria characterized factor four as "unquestionably the most important single factor" in the fair use determination, [5,6]and explicitly criticized Meta's strategy of attempting to use the strength of factor one to eclipse the analysis of factor four. The court for the first time in an AI training copyright case categorized the market harms assertable by plaintiffs into three sub-types, forming a layered analytical framework.

First, direct substitution: the model reproduces the original work or generates substantially similar content. The court found this path unsupported by the factual record in this case. According to expert testimony from both sides, even using specially designed adversarial prompts to induce Llama to recite training content, the model could not generate continuous fragments exceeding 50 words from any of the plaintiffs' books.[6]The court cited *Authors Guild v. Google, Inc.* as a reference point: in that case, Google Books allowed users to view snippets of up to 16% of books, yet the Second Circuit still found no significant harm to copyright value,

and Llama's output capability was far below this threshold—the causal chain for direct substitution therefore could not be established.[9,6]

Second, licensing market harm: the diminution of interest in the AI training licensing market. The court found this path not legally assertable, with its legal foundation lying in the requirement of systemic logical coherence. The court pointed out that if the "licensing market" were defined to include the potential licensing fees for the very use being complained of, then the plaintiffs' loss in that market would be a logical inevitability—from which it would follow that any unauthorized use would automatically constitute market harm, factor four would circularly favor the copyright holder in every case, thereby nullifying the entire normative effect of the fair use doctrine.[6]The court cited *Tresóna Multimedia, LLC v. Burbank High School Vocal Music Association*, *Bill Graham Archives v. Dorling Kindersley Ltd.*, and *Oracle* for the established principle: licensing fee losses resulting from transformative use do not constitute cognizable harm under factor four.[15,1,3,6]It should be noted here that the court did not deny the objective existence or possible development of an AI training licensing market as a factual matter; what it denied was the propriety of including such licensing fee losses in the balancing under the normative analytical framework of fair use—this is a normative judgment aimed at maintaining systemic coherence, not an empirical description of whether such a market exists.

Third, market dilution: the systematic construction of indirect substitution effects. This is the theoretically most innovative analytical link in this case's ruling, and the core position distinguishing Judge Chhabria from Judge Alsup. The court characterized this path as having "the most potential," but as unsupported in this case due to the plaintiffs' failure to build a sufficient evidentiary record. Its core proposition can be summarized as follows: even if an LLM can neither reproduce the original work nor generate substantially similar content, if it can produce, at extremely low marginal cost, large-scale expressive content similar to the original work in theme, style, or subject matter, thereby forming indirect market competition, then this "market dilution" effect likewise falls within the scope of factor four's concern for "the effect of the use upon the potential market for or value of the copyrighted work.[6]

A notable feature of this case's ruling on the "market dilution issue" is Judge Chhabria's open criticism of the core argument in the *Bartz* case, with this divergence centrally reflecting the deep methodological tensions in the current federal courts in the field of AI copyright. Judge Alsup in *Bartz* analogized the market competition effects generated by AI training to the competitive increment caused by "teaching children to write with books," and accordingly found that such effects were "not the kind of competitive or creative substitution that copyright law is concerned about." [21] Judge Chhabria directly rebutted this: "from the perspective of market effects, using books to teach children to write is far from creating a product with which a single individual can produce competitive works in bulk with minimal time and creativity—this improper analogy cannot be a basis for ignoring the most important factor in fair use analysis." [6] The fundamental divergence between the two judges is not in their judgment of transformative use itself—both cases found LLM training to be highly transformative—but in their understanding of the relationship between transformative use and market harm: Judge Alsup tended to regard high transformativeness as an almost decisive factor excluding adverse findings on factor four, a

stance reflecting a judicial orientation favoring the development of AI technology; Judge Chhabria insisted that transformative use and substantial market harm can coexist, and that even if the purpose of use is highly transformative, if the market dilution effect it triggers is sufficient to "seriously undermine the incentive for human creation," fair use should still be denied[6]The institutional significance of this divergence extends far beyond the individual cases: it essentially touches upon the issue of weight allocation among the four factors of fair use—whether the transformativeness analysis of factor one occupies a dominant position, or whether the market impact analysis of factor four imposes a constraining limitation on the former.

Surveying Thomson Reuters, Bartz, and Kadrey together, the directness versus indirectness of the competitive relationship constitutes the key variable causing the divergence in adjudicative outcomes. In Thomson Reuters, the involved AI was a non-generative legal search engine, whose purpose of use functionally coincided directly with the original work market, making market substitution self-evident, and the fair use defense therefore inevitably failed.[14] Bartz and Kadrey both involved general-purpose generative LLMs, making the competitive relationship indirect, but the fundamental divergence between the two judges on the justiciability of indirect competition determined the value orientation of their respective decisions. Judge Chhabria, at the end of his ruling, limited the precedential effect of this case with rare restrained language: "this ruling does not represent that Meta's use of copyrighted materials to train its language models is lawful. It merely represents that these plaintiffs made the wrong arguments, and failed to build an adequate record on the right ones." [6]The judicial signal conveyed by this statement is clear and far-reaching: copyright law's regulation of AI training will not automatically retreat due to technological breakthroughs; the construction of market dilution theory, although not applied in this case, has prepared for future plaintiffs a litigation path endorsed affirmatively by the judiciary, the reach of which will far exceed the interests of the individual parties in this case.

CONCLUSION

Overall, AI training cannot presently be characterized as categorically protected by fair use or categorically prohibited by copyright law. Future judicial decisions are likely to continue examining the purpose and transformative character of the use, the nature and amount of the copyrighted material, the legality of the source of the training data, and the effect on relevant markets. These factors will remain important in determining the balance between copyright protection and the development of artificial intelligence.

The cases also demonstrate the importance of the way training data is obtained and the possible effects on existing and future markets. Bartz highlights that the use of unlawfully obtained copies may be treated differently from the use of lawfully acquired materials, while Kadrey shows the importance of evidence concerning actual or potential market dilution.

A central consideration is the functional relationship between the AI system and the original works. Where the AI product has a direct competitive relationship with the original work, as discussed in Thomson Reuters, the fair use claim is more difficult to establish. In contrast, where

the AI system serves a different functional purpose, as considered in Bartz and Kadrey, the transformative-use analysis may provide stronger support for fair use.

The analysis of Thomson Reuters v. ROSS Intelligence, Bartz v. Anthropic, and Kadrey v. Meta shows that the application of the fair use doctrine to AI training data depends on the specific facts and technological context of each case. The courts have not established a single rule that makes the use of copyrighted works for AI training automatically fair or unfair.

Article Publication Details

This article is published in the [OpenMind Journal of Multidisciplinary Innovation & Development](#), ISSN: 3108-2920 (Online). In Volume 2 (2026), Issue 5 (Sep - Oct) - 2026
The journal is published and managed by [OMR PUBLICATION](#) .

Acknowledgements

We sincerely thank the editors and the reviewers for their valuable suggestions on this paper.

Authors' contributions

The author read and approved the final manuscript.

Funding

The author declare that no funding was received for this work.

Data availability

No datasets were generated or analyzed during the current study.

Declarations

Ethics approval and consent to participate

The author declares that not applicable.

Consent for publication

The author declares that not applicable.

Competing interests

The author declares no competing interests.

REFERENCES

1. Bill Graham Archives v. Dorling Kindersley Ltd., 448 F.3d 605 (2d Cir. 2006).
2. Campbell v. Acuff-Rose Music, Inc., 510 U.S. 569 (1994).
3. Google LLC v. Oracle America, Inc., 593 U.S. 1 (2021).
4. Hachette Book Group, Inc. v. Internet Archive, 115 F.4th 163 (2d Cir. 2024).
5. Harper & Row, Publishers, Inc. v. Nation Enterprises, 471 U.S. 539 (1985).
6. Kadrey v. Meta Platforms, Inc., 788 F. Supp. 3d 1026 (N.D. Cal. 2025).
7. Kelly v. Arriba Soft Corp., 336 F.3d 811 (9th Cir. 2003).
8. Authors Guild, Inc. v. HathiTrust, 755 F.3d 87 (2d Cir. 2014).
9. Authors Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015).
10. Sony Computer Entertainment, Inc. v. Connectix Corp., 203 F.3d 596 (9th Cir. 2000).
11. Sony Corp. of America v. Universal City Studios, Inc., 464 U.S. 417 (1984).
12. Shutterstock. (2023, July 11). Shutterstock partners with OpenAI and expands generative AI offerings [Press release].
13. Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc., 694 F. Supp. 3d 467 (D. Del. 2023).
14. Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc., 765 F. Supp. 3d 382 (D. Del. 2025).
15. Tresóna Multimedia, LLC v. Burbank High School Vocal Music Association, 953 F.3d 638 (9th Cir. 2020).
16. United States Code. (2018). 17 U.S.C. § 102(b).
17. United States Code. (2018). 17 U.S.C. § 107.
18. Warhol Foundation for the Visual Arts, Inc. v. Goldsmith, 598 U.S. 508 (2023).
19. White v. West Publishing Corp., 29 F. Supp. 3d 396 (S.D.N.Y. 2014).
20. Sega Enterprises Ltd. v. Accolade, Inc., 977 F.2d 1510 (9th Cir. 1992).

21. Bartz v. Anthropic PBC, 787 F. Supp. 3d 1007 (N.D. Cal. 2025).

PUBLISHER'S NOTE

OMR PUBLICATION remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of OMR PUBLICATION and/or the editor(s). OMR PUBLICATION disclaims responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.