Available Online: 4 November 2026 DOI: 10.54254/2753-7102/2026.38668

Examining artificial intelligence training as fair use: a comparative law reflection on Bartz v. Anthropic PBC

Dongyang Li[1,2]

[1]Faculty of Law, Macau University of Science and Technology, Macao, China
[2]Macao Intellectual Property Research Centre, Macau University of Science and Technology, Macao, China

wymtyyds@gmail.com

Abstract. Traditional legal doctrinal analysis easily regards the use of copyrighted works in Artificial Intelligence (AI) training as infringement; however, in substance this conclusion goes against the purpose of copyright law to balance interests. In Bartz v. Anthropic PBC, a U.S. federal district court held for the first time that the use of copyrighted books to train large language models is fair use, but did not accept fair use for the pirated copies acquired to build a permanent central library. Transformative use forms the core of this holding; training aims to build a technical tool, its results are transformative and differ in value from the original works, and no market substitution follows. The class action settlement reached later only concerned the acquisition of pirated copies; thus, the reasoning still stands. From the perspective of comparative law and legal interpretation, the decision offers a reference for China; Article 24 of the Copyright Law of the People's Republic of China lists fair use in a closed enumeration. It is better to give this provision a limited expansive interpretation. By means of teleological expansion, AI training can be included in the scope of "illustrating a point" under Article 24(1)(2) and treated by analogy as a special kind of illustrative use; its object is the linguistic patterns, statistical regularities, and knowledge structures in data, and three restrictive requirements apply. Then the degree of transformation serves as an auxiliary criterion and works inside the three-step test as a residual check.

Keywords: artificial intelligence training, fair use, transformative use, comparative law, legal interpretation

1. Introduction

Artificial intelligence technology keeps developing, and its applications now cover almost every area of social life; step by step, the material and cultural lives of the public have become richer. At the same time, the application of artificial intelligence has raised many legal questions, especially copyright questions, and scholars have debated them a lot in recent years. Scholars have carried out in-depth research on whether content generated by artificial intelligence is a work and whether artificial intelligence can be the subject of copyright [1-4]. With the deepening of these discussions, the use of other people's works during Artificial Intelligence (AI) training has also become a disputed issue in copyright law, and no answer has been reached so far.

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

AI training means a process in which AI obtains intelligence through preset programs and forms a model that can make decisions or carry out tasks on its own [5]. In this process, web crawler technology collects data on the internet, and deep learning and other related methods are used to learn semantics, and innovative output is the result. Others' works may be part of the training data, and at present one of the most debated questions is whether using them infringes copyright or counts as fair use under *the Copyright Law of the People's Republic of China* (hereinafter *the Copyright Law*). Some scholars have put forward a special copyright exception for text and data mining to regulate the use of others' works in AI training and to make the lawful boundaries of such use clear [6]. Other scholars suggest placing AI training under the statutory license [7]. Many scholars believe that AI training does not harm the value of works or the legitimate interests of copyright owners; therefore, they hold that it should be treated as fair use and regulated by specific rules [8-10]. These proposals have theoretical value and legitimacy, but legislation needs a long time and its result is highly uncertain; thus, legislation often cannot meet in time the judicial needs brought by technological development. Compared with legislation, the interpretive approach has more practical operability.

In the world, there are still few judicial decisions on the use of others' works in AI training, and most jurisdictions have been cautious about this question[1]. However, on June 23, 2025, the U.S. District Court for the Northern District of California held in *Bartz v. Anthropic PBC* (hereinafter *Bartz*) that the use of others' works in AI training is fair use, but it did not find fair use for the pirated copies kept in a central library[2]; this determination has great significance for judicial practice around the world. A close reading of the decision helps to answer how AI training should be assessed as fair use from the perspective of Chinese law. The case later ended in a class action settlement addressing the acquisition of pirated copies that the decision had left for trial, and the court granted final approval of the settlement and dismissed the entire action on July 20, 2026[3].

2. The interpretive dilemma of AI training under copyright law

2.1. AI training under traditional copyright theory

It is generally understood that AI training relies on deep learning models, and deep learning is in essence a process of learning from data, in which massive amounts of data are analyzed to "extract important patterns and trends" [11]. Technically, AI training uses data in two key stages, the data acquisition stage and the model training stage. In the data acquisition stage, the typical technical means by which AI obtains training data is collection through web crawler technology. In this process, websites usually establish a technical channel of communication with crawler programs through the Robots Exclusion Protocol, which specifies the scope and domains of web page data that may or may not be crawled [12]. In the model training stage, AI makes use of the crawled data. Natural language processing techniques convert the crawled text into machine-readable data, a model is built after screening and annotation, and the model parameters are optimized through iterative computation. In short, this stage is a process in which AI builds models by learning from data, performing only a technical analysis of the expression in works without extracting their underlying ideas [13].

---

1 See Thomson Reuters Enterprise Centre GmbH v. ROSS Intelligence Inc., 765 F. Supp. 3d 382 (D. Del. 2025), aff'd, No. 25-2153 (3d Cir. Sept. 29, 2026). See also Getty Images (US) Inc. v. Stability AI Ltd., [2025] EWHC 2863 (Ch) (Eng.).

2 See Bartz v. Anthropic PBC, 787 F. Supp. 3d 1007 (N.D. Cal. 2025) (order on fair use).

3 See Bartz v. Anthropic PBC, No. 4:24-cv-05417-AMO, ECF No. 680 (N.D. Cal. July 20, 2026) (order granting final approval of class action settlement).

The data acquisition stage thus supplies data for AI training, whereas the model training stage is the core of AI training, its goal being, in the words of one textbook, to develop "methods and algorithms that 'learn' from data" [14], so as to achieve specific purposes in subsequent applications.

Copyright scholars' interest in AI training centers on its crawling and use of data. Where the data used for AI training consist of copyrighted works, such crawling and use may infringe copyright. Crawling the data of others' works in AI training resembles reproduction in the copyright sense, yet this reproduction does not exist in isolation and instead constitutes incidental copying or intermediate copying [15]. Once copied, the information is used to construct an algorithmic model, so the nature of the act is closely tied to its ultimate use. Although the act is intermediate and functional at the technical level, at the legal level it may still infringe the right of reproduction in others' works. Traditional copyright theory holds that, whether the reproduction can be directly perceived or occurs through intermediate machine reading, the elements of reproduction are satisfied as long as the act produces "one or more copies" of a particular work. Accordingly, even if data crawling and preprocessing are merely technical steps in the training process, the comprehensive reproduction of works without permission should still be regulated by copyright law, and reproduction is the central issue in AI training [16].

The use of work data in the model training stage may also infringe the right of adaptation and other rights in others' works [17-19]. As noted above, although AI systems aim to extract the features of an algorithmic model by learning from data rather than simply presenting the content of works to users, their technical processing of protected original expression during training does in a sense alter the original form of expression of the works, giving rise to a kind of transformative use that serves machine learning. Absent further legal provisions or express authorization by the right holder, however, whether such transformation constitutes fair use or falls within the scope of the right of adaptation remains highly contested in both scholarship and practice.

2.2. The interpretive dilemma: doctrinal interpretation and the loss of substantive reasonableness

Article 10 of *the Copyright Law* lists the exclusive rights of copyright owners by type; these rights are exclusive, so other people cannot carry out the restricted acts without the permission of the right holder, and otherwise infringement arises in principle. In addition, Article 24 of *the Copyright Law* uses a closed (exhaustive) enumeration, and with it the law expressly limits the specific situations of fair use that are exempt from infringement[4]. Under strict legal doctrinal analysis, right holders have not authorized AI training and no statutory exception expressly covers such training; therefore, in formal logic, a finding of copyright infringement is easy to reach. On the other hand, doubts exist about whether this conclusion is reasonable in substance.

2.2.1. AI training does not substantially harm the legal interests protected by copyright law

The types of conduct that lead to liability for copyright infringement are listed in detail in Articles 52 and 53 of *the Copyright Law*. However, meeting the constituent elements of the conduct is not enough for a finding of infringement; its substantive core must be harm to the legitimate rights and interests of the copyright owner. This value orientation is in line with the spirit of the three-step test set up by the Berne Convention; under this test, rights may be limited in certain cases on the condition that the limitation does not conflict with the normal use of the work or unreasonably harm the author's legitimate interests.

[4] See Copyright Law of the People's Republic of China (as amended by the Standing Committee of the National People's Congress on November 11, 2020) art. 24, http://en.npc.gov.cn.cdurl.cn/2020-11/11/c_674691_2.htm.

AI training aims to extract statistical regularities and pattern features from texts by algorithmic analysis and data learning, and in the end it forms predictive models that can be used for downstream tasks; its aim is not to reproduce or disseminate the expressive content of the original works. In technical terms, it is more like reading and analysis, and less like reproduction and distribution. Such training does not replace the market demand for the original works, and it does not harm their communicative value or the income of authors. Training uses a huge number of works in a highly dispersed and non-expressive way, so training can almost never compete with and replace the market of any particular work [20]. Therefore, the use does not cause the substantial harm that copyright law aims to prevent, and to treat it directly as infringement goes against the legal philosophy of copyright law, whose center is the balance of interests. Even if the generated content is later found to infringe, the data processing in training should not be regarded as infringing after the fact.

2.2.2. Training accords with the value orientation and functional role of the fair use system

Fair use is an important balancing mechanism in the copyright system, and it puts necessary limits on the exclusive rights of right holders in order to ensure that the public can reasonably access and use works; in this way, it promotes cultural progress and the spread of knowledge [21]. Article 24 of *the Copyright Law* does not expressly list AI training as one of the specific situations of fair use, but this does not rule out finding room for its legitimacy in the existing system through legal interpretation or gap filling.

From the point of view of institutional purpose, whether a use is fair use should not be decided by mechanically matching the form of conduct. Instead, the question is whether the use fits the incentive goals of copyright law, and also whether a certain limitation on the copyright owner helps to achieve a greater social benefit. AI training learns from a huge number of works in an efficient way, and as a result the level of machine intelligence has clearly risen and scientific research, technological innovation, and cultural development have advanced; the social benefits of training are far greater than its potential effects on copyright owners [22]. Most importantly, training uses works to build technical models and not for expressive purposes, so it is a typical non-expressive use. It does not compete with the market for the original works or affect the economic returns of copyright owners; therefore, it meets the core requirements of fair use in Article 24 of *the Copyright Law*, that is, "the normal use of the work is not affected and the legitimate rights and interests enjoyed by the copyright owner are not unreasonably prejudiced"[5].

Thus, it is hard to justify in substance the infringement conclusion that formal logic leads to. Recently, the U.S. District Court for the Northern District of California held for the first time that the use of copyrighted books to train large language models is fair use, and the ruling gives a comparative law reference as well as a path to reexamine how well traditional legal doctrinal analysis fits new technologies.

3. The characterization of AI training in U.S. judicial practice

3.1. The reasoning of the Bartz decision

Three American authors and their companies brought the copyright infringement action against Anthropic. According to the plaintiffs, Anthropic copied their copyrighted books without authorization and used them to train the Large Language Models (LLMs) behind its AI service Claude. The way Anthropic obtained and used the books was at the center of the case. Anthropic downloaded more than seven million copies of books from pirate sites in 2021 and 2022, and from 2024 onward it bought millions of print books, took off their bindings,

[5] Copyright Law of the People's Republic of China; see note 4.

and scanned them into digital form. In training the LLMs, these copies were cleaned, converted into tokens (the smallest semantic units of text that a model processes, a core concept of natural language processing in AI training), and compressed, and the aim was to improve the ability of Claude to generate text. The plaintiffs claimed that these acts of copying infringed their copyrights, and Anthropic raised the defense of fair use. The court held that "the use of the books at issue to train Claude and its precursors was exceedingly transformative and was a fair use", that the digitization of the print books purchased by Anthropic was also a fair use, and that "Anthropic had no entitlement to use pirated copies for its central library", because "[c]reating a permanent, general-purpose library was not itself a fair use excusing Anthropic's piracy"[6]. The legality of the data source thus bears on whether a fair use defense can succeed.

The court assessed three distinct uses by the defendant separately, structuring its analysis around the four fair use factors in Section 107 of the U.S. Copyright Act. As to the books obtained from pirate sites, the court declined to find fair use. It stated that piracy of copies that could have been purchased or otherwise accessed lawfully "is inherently, irredeemably infringing even if the pirated copies are immediately used for the transformative use and immediately discarded", while adding that it "need not decide this case on that rule". Stressing that "[a] separate justification was required for each use", the court rejected Anthropic's defense and left liability and damages for the pirated copies to trial. As to LLM training and the digitization of print books, the court found fair use. On the purpose and character of the use, training LLMs to generate new content by learning the statistical relationships in text was "exceedingly transformative", while digitizing print books merely changed the storage format and involved no distribution outside the company. As for the nature of the copyrighted works, the court accepted that the second factor "points against fair use for all copies alike", but it considered that its main function was "to help assess the other factors". Regarding the amount used, entire books were copied, but the court accepted that "Anthropic needed billions of words to train any given LLM". With regard to market effect, users could not obtain infringing output due to technical measures; thus, the training did not cause substitution in the market for the original works[7]. The decision also reflects a rule, that is, the legality of the data source does not depend on the later use, and it sends a signal that companies should obtain training data through lawful channels.

In form, the court came to the conclusion that some of the uses were fair through the four-factor analysis, but its central reasoning has a close link with the concept of transformative use. The court found Anthropic's AI training to be fair use mainly on the ground that the training was a transformative use. The court stated that "[l]ike any reader aspiring to be a writer, Anthropic's LLMs trained upon works not to race ahead and replicate or supplant them" but "to turn a hard corner and create something different", and the statement puts into concrete words the technical view that model training builds a data learning environment where intelligent capabilities are formed. According to the express finding of the court, Anthropic used the works "to iteratively map statistical relationships between every text-fragment and every sequence of text-fragments" so that the finished LLM could produce new text, and the essence of the use was to "create something different"[8]. In substance, the conclusion expresses a central point, namely that training has a purpose fundamentally different from the creative purpose of the original works and does not replace their literary and artistic value.

To be more specific, in this case the transformative purpose of AI training has two levels, a technical one and a legal one. At the technical level, AI "understanding" is a cognitive mechanism formed by millions of

6 Bartz, 787 F. Supp. 3d 1007; see note 2.

7 Bartz, 787 F. Supp. 3d 1007; see note 2.

8 Bartz, 787 F. Supp. 3d 1007; see note 2.

features and the interactions among them. This is similar to how humans take in knowledge after reading; however, in technical terms it is realized by extracting linguistic regularities and logical structures with algorithms. The court pointed out that "Authors do not allege that any LLM output provided to users infringed upon Authors' works", and it also noted that there was "no allegation of any traceable connection between the Claude service's outputs and Authors' works" and that the Claude service "placed additional software between the user and the underlying LLM to ensure that no infringing output ever reached the users"[9]. It can be seen that training ends in building a technical tool, not in disseminating or reproducing the original works, and that it has a new functional character. At the legal level, the court compared training to the way people learn to write through reading and said that "[e]veryone reads texts, too, then writes new texts"; it also noted that, under Section 102(b) of the U.S. Copyright Act, copyright does not extend to "method[s] of operation, concept[s], [or] principle[s]"[10]. AI's extraction of linguistic regularities is exactly in this unprotected area, and it further confirms that, in legal terms, AI training goes beyond the original purpose of the works.

3.2. Subsequent developments: the settlement and its corroboration of the decision's rules

After the decision, the case ultimately ended in a class action settlement. On July 17, 2025, the court certified a class limited to the copyright owners of books in the versions of the Library Genesis (LibGen) and Pirate Library Mirror (PiLiMi) datasets downloaded by Anthropic. The books that Anthropic had purchased and scanned, as well as the training use, fell outside the class claims, since the June 2025 decision had already held both to be fair uses[11]. The parties signed a term sheet on August 25, 2025, and on September 5 they signed a settlement agreement and sought preliminary approval. The court granted preliminary approval at the close of a hearing on September 25 and issued a memorandum opinion on preliminary approval on October 17[12].

The settlement agreement provides for a non-reversionary settlement fund of $1.5 billion covering the 482,460 works on the Works List[13], which yields "a per-work award of about $3,000, less costs and fees"[14]. Anthropic must also destroy all original files of works downloaded from LibGen or PiLiMi, together with any copies originating from them, subject to certain legal preservation obligations[15]. The release is confined to past input claims. Class members give up claims "broadly related to past AI inputs, namely pirating and copying their works in the past up to the point of any AI output", but they do not give up "claims about past AI outputs, nor claims of any kind about future conduct (on or after August 25, 2025)"[16]. On July 20, 2026, the court granted final approval of the settlement, entered final judgment, dismissed the entire action with prejudice, and closed the case[17]. The two notices of appeal filed thereafter came from law firms that had been excluded from the fee award, and the parties reported that both related solely to the fee award. No class member or objector filed a timely appeal, the parties agreed that the effective date of the settlement had

---

9 Bartz, 787 F. Supp. 3d 1007; see note 2.

10 Copyright Act of 1976, 17 U.S.C. § 102(b).

11 See Bartz v. Anthropic PBC, 791 F. Supp. 3d 1038 (N.D. Cal. 2025) (order on class certification).

12 See Bartz v. Anthropic PBC, No. C 24-05417 WHA, 2025 WL 2961371 (N.D. Cal. Oct. 17, 2025) (memorandum opinion on preliminary approval of class action settlement).

13 Bartz, ECF No. 680; see note 3.

14 Bartz, 2025 WL 2961371; see note 12.

15 Bartz, ECF No. 680; see note 3.

16 Bartz, 2025 WL 2961371; see note 12.

17 Bartz, ECF No. 680; see note 3.

passed, and distributions to class members were expected to begin no later than November 15, 2026[18].

What the settlement resolved was precisely the acquisition of pirated copies that the June 2025 decision had left for trial. The decision's findings that the training use and the digitization of print books were fair uses were neither vacated nor reversed. Because the certified class was limited to books in the LibGen and PiLiMi datasets, the $1.5 billion settlement consideration was directed at liability and damages arising from the acquisition of pirated copies, while the training use had already been held to be fair use. The rule that the legality of the data source does not depend on the later use gets further support from this outcome. These fair use findings nonetheless come from a summary judgment ruling of a federal district court, were never subjected to appellate review because the case settled, and do not bind other courts.

3.3. The elaboration and implications of the decision's rules

Summarizing the spirit of the decision further, beyond treating transformative purpose as the primary criterion of fair use, the case offers other lessons, such as that market substitution must be proven, that the legality of the data source affects the path by which a fair use defense may succeed, and that the quantity of copying is not in itself a negative factor. These rules, as reflected in the decision, in substance represent the court's consideration of the degree of transformation reached by AI training that has a transformative purpose. The court's detailed analysis of parts of the four fair use factors can be summarized as an assessment of the degree of transformation of Anthropic's LLM training along two dimensions, the market substitution effect and the creation of new social value.

3.3.1. Market substitution denied: technology-tool markets distinguished from content markets

In its analysis of the fourth factor, the court stated expressly that "[t]he copies used to train specific LLMs did not and will not displace demand for copies of Authors' works, or not in the way that counts under the Copyright Act", which indicates that AI training targets the market for technical tools rather than the content creation market. Training LLMs "did not result in any exact copies nor even infringing knockoffs of their works being provided to the public". The court even took for granted that each LLM had "memorized" the works on which it was trained almost verbatim, yet filtering software at the output end ensured that infringing output never reached users. The court acknowledged that "Claude could help less capable writers create works as well-written as Authors' and competing in the same categories" and assumed that training LLMs would result in an explosion of works competing with the plaintiffs' works. It reasoned, however, that this complaint was no different from complaining that training schoolchildren to write well would result in an explosion of competing works, and that "[t]his is not the kind of competitive or creative displacement that concerns the Copyright Act", since "[t]he Act seeks to advance original works of authorship, not to protect authors against competition". The plaintiffs further argued that training LLMs had displaced, or would displace, an emerging market for licensing their works for the narrow purpose of training LLMs. Even assuming the plaintiffs to be correct, the court held that "such a market for that use is not one the Copyright Act entitles Authors to exploit"[19]. It may be inferred that copyright protection of potential markets is bounded by the interests foreseeable to the original author, which further indicates that the risk of AI training diminishing the market value of the original works is very low. The non-substitutive manner of Anthropic's training use and its confinement within the reasonable boundaries of potential markets together support the degree of the

[18] See Status Report, Bartz v. Anthropic PBC, No. 4:24-cv-05417-AMO, ECF No. 688 (N.D. Cal. Sept. 2, 2026).

[19] Bartz, 787 F. Supp. 3d 1007; see note 2.

reasonable boundaries of potential markets together support the degree of transformation of the AI training at issue in market and economic terms.

3.3.2. Extending social value: technological empowerment and enhanced public welfare

The decision implicitly recognizes the social value of AI technology. Although the court did not directly address the social value of the defendant's AI training, it regarded the technology at issue as "among the most transformative" that many people will see in their lifetimes[20]. The potential of this technology to raise productivity and improve public services is highly consistent with the legislative purpose of fair use, namely promoting the progress of science and useful arts[21], which illustrates the positive externalities of technological innovation. AI researchers have observed that high-quality training data is a precondition for the generalization ability of AI (telephone interview conducted by the author, 7 February 2026). The court's finding that the digitization of print books was a fair use in effect provides companies with a compliant path to obtaining data through lawful means and implies an expectation that fair use will promote the optimal allocation of data resources. In the court's reasoning, this use had the transformative purpose of changing the storage format, since every purchased print copy "was copied in order to save storage space and to enable searchability as a digital copy", and such conversion "was transformative for that reason alone", which led the court to find a transformative use[22]. From the perspective of social value, this finding can be read as recognizing that the use had reached a considerable degree of transformation. The outcome of the decision makes clear the court's recognition of the degree of transformation of AI training in creating new social value.

The concept of transformative use recurs throughout the lengthy decision and is a key pillar on which the court based its fair use ruling. It therefore holds considerable value for resolving the interpretive dilemma that Chinese copyright law faces regarding the use of training data in AI training. In China, however, transformative use remains at the stage of a scholarly concept. For it to operate in the context of Chinese law, the concept requires systematic elaboration and adaptation to the Chinese legal system.

4. Transformative use: the legal essence of AI training

4.1. The nature and elements of transformative use

Transformative use refers to a secondary use that "adds value to the original" [23]. It requires that "a secondary work add new content on the basis of the original, with a further purpose or a different character, creating new information, new aesthetics, and new insights and understandings" [24]. Empirical research on a large body of U.S. case law has found that transformative use is "a robust predictor of a finding of fair use" and has pointed to "the dominance of transformative use over other factors" [25].

The doctrine originated in the 1994 decision of the U.S. Supreme Court in *Campbell v. Acuff-Rose Music, Inc.*, which sought to resolve the vagueness surrounding the purpose and character of the use, one of the four traditional fair use factors[23]. Subsequent decisions in *Blanch v. Koons* and *Seltzer v. Green Day, Inc.* further consolidated the judicial standing of transformative use. The former case established that a shift of purpose

[20] Bartz, 787 F. Supp. 3d 1007; see note 2.

[21] See U.S. Const. art. I, § 8, cl. 8.

[22] Bartz, 787 F. Supp. 3d 1007; see note 2.

[23] See Campbell v. Acuff-Rose Music, Inc., 510 U.S. 569 (1994).

from aesthetic experience to social commentary can be transformative use[24]; the latter case made it clear that transformative use also covers a later work that brings a wholly new aesthetic experience and informational value through creative transformation[25]. The above decisions follow the same logic. If a later work has a different purpose from the original work and creates substantial new value, it may be considered fair use.

U.S. judicial practice also divides transformative use into two types, content transformation and functional transformation [26], and the two may exist separately or apply together. (The transformations in the cases above are all content transformations.) In *Bill Graham Archives v. Dorling Kindersley Ltd.*, the court held that the defendant's copying constituted transformative use because it served the new functional purpose of compiling a historical record, thereby confirming the independent value of functional transformation[26]. Although content transformation and functional transformation differ, both point to the core principle for establishing transformative use, namely that the user subjectively has a transformative purpose and objectively produces a transformative result.

In the traditional fair use analysis, a commercial purpose usually weighs against a finding of fair use. The emergence of the transformative use doctrine broke through this rigid view, since "as long as the use has a purpose different from that of the original and produces a transformative result, the commercial value of the later work does not necessarily preclude the application of fair use" [27]. This shift in judicial thinking has made the combination of a new purpose or function with a transformative result the core yardstick for determining whether transformation exists.

The existence of transformation does not necessarily lead to a finding of fair use. Its degree must reach a legal threshold, that is, a considerable degree of transformation. U.S. courts have long had to gauge the degree of transformation, and several possible yardsticks can be identified, such as the proportion of the new work taken from the original and the time invested in the secondary creation, as well as the extent of market impact considered in the case law[27]. Among these, the creation-time criterion lacks general applicability because works differ by type. Given that copyright law functions as a tool of market incentives, the analysis of the degree of transformation should focus on economic impact.

The market failure theory proposed by one commentator [28] provides a theoretical foundation for analyzing the degree of transformation. Its core view is that the application of fair use must satisfy two conditions, namely that the use does not substantially harm the copyright owner's incentive to create and that it promotes the public interest [29]. This accords with the legislative purpose of copyright law. The concept of transformative use was created precisely to address the new ways of using works brought about by technological development, promoting the dissemination of knowledge and cultural innovation while safeguarding the rights of authors. Economic considerations should therefore focus on two core dimensions, the substitution effect of the later work on the market position of the original and the degree to which new social value is created.

Through long judicial exploration, U.S. courts have developed a relatively mature chain of reasoning for transformative use. They first examine whether the use has a transformative purpose or function and produces a transformative result. They then measure the degree of transformation by considering the substitution effect of the later work on the market position of the original and the degree to which new social value is created.

---

[24] See Blanch v. Koons, 467 F.3d 244 (2d Cir. 2006).

[25] See Seltzer v. Green Day, Inc., 725 F.3d 1170 (9th Cir. 2013).

[26] See Bill Graham Archives v. Dorling Kindersley Ltd., 448 F.3d 605 (2d Cir. 2006).

[27] See Ty, Inc. v. Publications International Ltd., 292 F.3d 512 (7th Cir. 2002).

Finally, they assess the likelihood that fair use is established according to the degree of transformation. This progressive framework provides an important reference for resolving copyright disputes in the age of AI.

4.2. AI training as transformative use

4.2.1. AI training has a transformative purpose and produces transformative results

In the data acquisition and learning stage, the purpose for which model training uses works differs fundamentally from the original creative purpose of the prior works. Model training is in essence a process of building a data learning environment in which AI, by analyzing the regularities and features underlying data, gradually forms intelligent capabilities and ultimately achieves preset task objectives. The understanding of AI is in essence a cognitive mechanism formed through millions of features and billions of interactions among them, and such "understanding" of the basic linguistic elements underlying human expression and creation does not fall within the subject matter protected by copyright law. In short, use at this stage goes beyond the original purpose of the works and carries the attribute of a new purpose.

In the output stage of AI training, the core objective is not to copy prior works but to generate technical results for a wider range of applications. This shows that the use of works in model training has the character of a new function. Although the process is usually commercial, it does not involve appreciative use of the aesthetic value of the original works in literature, art, or science. Instead, technical processing yields a transformative result in the form of a pattern-based model. This result does not directly conflict with the original works in form of expression, functional positioning, or field of application, and it fully meets the criteria for transformation summarized above.

4.2.2. The Degree of Transformation: AI Training Does Not Affect the Market Position of the Original Works

One of the core factors in determining fair use is the analysis of market substitution, that is, whether the use at issue substantially harms the market value of the original works [30]. Viewed from the training purpose of large AI models, their marked transformative character is an important premise for assessing market impact. Under the rule laid down by the U.S. Supreme Court in *Campbell*, the degree of transformation stands in a marked inverse relationship to the risk of diminishing the market value of the original work[28]. In terms of market competition, although the intensive training of generative AI models requires massive and varied training data such as books and image materials, academic papers, and literary works, the content it generates is directed at emerging applications such as government information services, learning assistance, and intelligent customer service. This technical character means that AI training targets the market for technical tools rather than the content creation market in direct competition with the original works. In addition, AI-generated content is nondeterministic and unpredictable, and its technical mechanism differs entirely from the deterministic retrieval of the traditional model of memory storage and information reproduction [31]. In the traditional framework for analyzing market impact under copyright law, the key question is whether conduct carrying a risk of infringement would hollow out the right holder's legitimate interests in the work by presenting the work directly to third parties. In the context of AI training, the focus is on distinguishing direct copying from transformative use. Copying in AI training serves only the algorithm's deep learning of the linguistic regularities and patterns of thought implicit in the data. Such copies are temporary, existing only during the model training stage. Once the model has been produced it no longer depends on the original data, still less does it provide copies of the original works to third parties.

[28] Campbell, 510 U.S. 569; see note 23.

Moreover, the outputs of AI training and the original works usually belong to different fields of application, and their user groups are clearly separated, so no market effect arises that conflicts with the reproduction of the value of the original works. At the same time, the protection that copyright law affords to the potential markets of original works does not extend indefinitely but should be bounded by the interests that the author could reasonably foresee at the time of creation. Even if AI-generated content partially draws on the value of the original works in emerging fields of application, this does not affect the right holder's exclusive interests or creative incentives in traditional markets. In the industrial ecosystem, AI training and the production of original content by authors complement each other and do not substitute for each other. Therefore, with the purpose of use, the manner of conduct, and the output all taken into account, AI training has a low risk of reducing the copyright market value of the original works, and at this level its highly transformative character can be confirmed in a preliminary way.

4.2.3. The degree of transformation: AI training generates new social value

Fair use is a mechanism that keeps a dynamic balance between the protection of private rights and public welfare, and it also has the mission of pushing forward cultural innovation and the evolution of human civilization; both are part of its core value. Digital and intelligent civilization is developing faster and faster, and large AI models now create economic and social value that goes beyond the interests of certain developers or companies. These models raise total factor productivity, make social governance more effective, and improve the quality of public services; such spillover effects are clearly pushing up the overall capacity for innovation and the level of public welfare in society in a systemic way. The above positive externalities of technological innovation agree closely with the aim of fair use to maximize social value, and they give data processing in the training of large AI models an institutional basis of legitimacy. In addition, large AI models need a supply of high-quality, large-scale training data to capture human cultural characteristics, linguistic regularities, behavioral patterns, and knowledge systems accurately, and such supply directly decides how well a model generalizes and how widely it can be applied [32]. Due to this inherent need for open data, the key value of fair use in the optimal allocation of data resources becomes even clearer, and when the scope of exclusive copyright is moderately narrowed, the system has room for data to be used in non-commercial and transformative ways. Algorithmic bias in AI is a growing problem, and fair use can also do some good for technology ethics here by encouraging the building of more diverse training datasets. Developers may use works from many different sources in a compliant way, and then the training corpora can be more representative and data bias may cause less algorithmic discrimination; thus, besides its traditional function of protecting rights, copyright law also gains a regulatory value in promoting technological fairness [33].

Prior works have social value mainly because they provide aesthetic experience and pass on knowledge; this value encourages creators to carry out diverse intellectual creation and enrich original expression, and in the end it serves the institutional goal of making culture prosper and promoting innovation in society. As mentioned above, AI training has an obvious transformative character and functional novelty, and its technological applications have already gone beyond the traditional scope in which works are used. AI systems now show their potential for technological empowerment in many fields, such as conversational commercial interaction, the intelligent identification of risks in mines, natural language processing, and multimodal content generation. Traditional works in literature and art have appreciative value; by comparison, the social value of AI training is more about raising efficiency in real-world applications, and through the innovative use of technical tools it brings new incremental value to the overall interests of society. Such a creative extension of social value is further proof that AI training is highly transformative.

5. A limited expansive interpretation of AI training as fair use under Chinese law

The U.S. doctrine of transformative use comes from judge-made law; therefore, Chinese courts cannot use it directly as a basis for adjudication. In addition, Article 24 of *the Copyright Law* uses a closed enumeration, and neither AI training nor technical use is expressly included in it. Thus, conduct that is legitimate in substance does not easily fit the form of the current law. The solution should first be to update the normative content of the law in the current legal framework, and not to start over by creating a new category of fair use. This limited expansive path in the courts keeps the stability that *the Copyright Law* needs as a statute, and at the same time it can meet the special needs of the AI era through techniques like teleological expansion. In its basic position, the interpretive approach to AI training as fair use does not reject the closed nature of the categories in the current law. Instead, it uses methods of legal interpretation to place AI training in the textual scope or the normative purpose of the existing provisions.

5.1. Interpreting the transformative purpose of AI training under Chinese law

To bring AI training under Article 24(1)(2) of *the Copyright Law*, literal interpretation alone obviously does not work. This provision permits "appropriate quotation from a published work of another in one's own work for the purpose of introducing or commenting a certain work, or illustrating a point"[29], and the core of its text is the expressive use of a prior work, in which another's work is used to argue for or explain a specific point. AI's technical analysis is non-expressive and aims to extract linguistic regularities and statistical features, so the two are somewhat far apart. Therefore, several methods of legal interpretation have to be used together to close the gap between the statutory text and new forms of use.

From the angle of teleological interpretation, Article 24 aims to keep the basic balance in copyright law between encouraging creation and promoting the public interest, and its legislative purpose is not just to list some exemptions. AI training has great social value for technological innovation, the spread of knowledge, and public welfare. Interpreters might follow the closed text rigidly and keep AI training completely out of fair use; in that case there would surely be a logical break between the legislative purpose and the result, and copyright law should promote cultural progress but would then block the data use that one of the most transformative technologies needs. Therefore, interpreters should expand "illustrating a point" in a moderate teleological way. The point to illustrate need not stay within literary, artistic, or social issues, and it should also include technical questions, for example, the discovery and verification of linguistic regularities, logical structures, and algorithmic features. This analogy is legitimate due to functional similarity. Traditional scholarly quotation uses others' works to explain theoretical propositions, and in a similar way AI systems find and verify linguistic rules by analyzing works with algorithms; in both cases, works are a means to serve a purpose that is independent of their expressive value.

Further support comes from systematic interpretation. The twelve specific situations listed in Article 24 are not alike in the manner of conduct, and their common point is the external nature of the purpose of use. In these situations, the work is used mainly to serve other legitimate ends, such as education, scientific research, and news reporting, and not to reproduce its aesthetic or informational value. AI training uses works to build intelligent models, not to spread their content; in functional structure, this purpose is very similar to Article 24(1)(6) on scientific research and Article 24(1)(3) on news reporting. Thus, based on systematic

[29] Copyright Law of the People's Republic of China; see note 4.

interpretation, AI training can be treated as a new type of use with an external purpose, and the boundaries of the text are not exceeded.

In terms of historical interpretation, Article 24(1)(13) was added in the third revision of *the Copyright Law*, and it covers other situations provided for by laws and administrative regulations. This catch-all clause cannot be a direct basis for adjudication on its own; however, the fact that it exists shows that the legislature wanted to keep the fair use system open to some extent, and the legislature did not expressly forbid including technical uses in fair use.

However, strict institutional controls must go along with the use of analogy, so that overextension and abuse do not erode the closed character of fair use. For this reason, three restrictive requirements are needed, and they also answer and absorb the reasoning in *Bartz*. First, the source must be lawful. Training data must come through lawful channels; for work data that come from pirate websites or are obtained by other infringing means, a fair use defense is not available. Second, the use must be non-expressive. The aim of AI training cannot be to reproduce the original expression of works, and at the output end, technical measures should make sure that the generated content is not substantially similar to the original works. Third, there must be no market substitution. In existing or potential markets, training must not compete with and replace the economic value of the original works. AI training can be recognized as fair use along the path of limited analogy only on the condition that all three requirements are met; thus, a careful balance is kept between the rigidity of norms and the response to technology.

5.2. Integrating the degree of transformation of AI training into Chinese law

The highly transformative character of AI training has been fully shown at the empirical level and at the theoretical level, and an inherent correlation can be seen: the higher the degree of transformation, the lower the risk that the market for the original works is replaced. This correlation also gives the interpretive approach a substantive basis of legitimacy. However, it must be made clear that the degree of transformation is only related to whether a use can be justified, and it does not equal fair use in law. An interpretive bridge must be built to move from the comparative law theory of transformative use to its normative application in Chinese law, so that a scholarly concept can be lawfully turned into a basis for adjudication.

Therefore, when AI training is interpreted in law as fair use, the degree of transformation should be an auxiliary criterion, not a ground of decision that stands apart from the statutory requirements. The logic of its operation has three levels, one after another. First, teleological expansion is used to decide whether AI training fits the normative purpose of illustrating a point in function; in other words, the question is whether it uses works mainly to reveal and verify technical questions, such as linguistic regularities, knowledge structures, or algorithmic features, and not to reproduce their literary or artistic expression. Second, if the answer to the first question is yes, the degree of transformation is assessed with the analytical framework above, and the focus is on whether the use creates new social value and whether its effect on the market for the original works is substantially non-substitutive. Lastly, the three-step test works as a residual constraint, and it is used to review carefully whether the use affects the normal use of the original works and whether it unreasonably harms the legitimate rights and interests of the copyright owner. AI training can finally be found to be fair use only after it passes all three tests. Thus, the degree of transformation is considered inside the interpretive framework of the existing provision.

6. Conclusion

In Chinese law, there is no need to break through the current legislative framework and start again in order to recognize AI training as fair use. With various techniques of legal interpretation used together, AI training can fall within the normative scope of "illustrating a point" under Article 24(1)(2) of *the Copyright Law*. This limited expansive path through the courts does not bring the institutional risk of judge-made law, and it also gives the AI industry the legal flexibility it needs to use data lawfully. Of course, the interpretive approach also has limits. If AI training starts to use works in new ways, or if the risk of market substitution rises significantly, turning to the legislative approach again would still be a necessary option. At present, however, the interpretive approach is more prudent and pragmatic, and it is also more in line with the rule of law.

References

[1] Huang, H., & Huang, J. (2019). The reasonability of artificial intelligence products regarded as works [in Chinese]. *Jiangxi Social Sciences, 39*(2), 33–42, 254.

[2] Wang, Q. (2023). The qualitative analysis of content generated by artificial intelligence in copyright law [in Chinese]. *Tribune of Political Science and Law, 41*(4), 16–33. https://doi.org/10.3969/j.issn.1000-0208.2023.04.002

[3] Xu, C., & Yuan, Y. (2019). The legal subjectivity of artificial intelligence: From the perspective of copyright protection of AI products [in Chinese]. *Science Technology and Law, 2019*(2), 1–6, 18. https://doi.org/10.19685/j.cnki.cn11-2922/n.2019.02.001

[4] Yi, L., & Yin, S. (2021). The negation of the artificial intelligence's qualification of legal subject in copyright law of the People's Republic of China [in Chinese]. *Science Technology and Law (Chinese-English Version), 2021*(4), 66–72, 148. https://doi.org/10.19685/j.cnki.cn11-2922/n.2021.04.008

[5] Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development, 3*(3), 210–229. https://doi.org/10.1147/rd.33.0210

[6] Song, Y. (2017). The copyright infringement immunity of text and data mining: From the perspective of the proposal for EU copyright directive amendment [in Chinese]. *Electronics Intellectual Property, 2017*(6), 42–51. https://doi.org/10.3969/j.issn.1004-9517.2017.06.004

[7] Liu, Y., & Wei, Y. (2019). The copyright infringement issue of machine learning and its solution [in Chinese]. *ECUPL Journal, 22*(2), 68–79. https://doi.org/10.3969/j.issn.1008-4622.2019.02.006

[8] Liu, Y. (2025). Fair use routes of copyright for artificial intelligence training works [in Chinese]. *Journal of Northeastern University (Social Science), 27*(1), 117–126. https://doi.org/10.15936/j.cnki.1008-3758.2025.01.013

[9] Xu, X., & Yang, Y. (2019). On the fair use of copyright in deep learning of artificial intelligence [in Chinese]. *SJTU Law Review, 2019*(3), 32–42. https://doi.org/10.19375/j.cnki.31-2075/d.2019.03.003

[10] Jiao, H. (2022). Copyright risks and solutions for data acquisition and utilization in artificial intelligence works [in Chinese]. *Contemporary Law Review, 36*(4), 128–140. https://doi.org/10.3969/j.issn.1003-4781.2022.04.012

[11] Hastie, T., Tibshirani, R., & Friedman, J. (2001). *The elements of statistical learning: Data mining, inference, and prediction*. Springer. https://doi.org/10.1007/978-0-387-21606-5

[12] Zhang, P. (2013). The general clause of the Anti-Unfair Competition Law and its application: Reflections prompted by search-engine crawler protocols [in Chinese]. *Journal of Law Application, 2013*(3), 46–51.

[13] Zhou, Z. (2016). *Machine learning* [in Chinese]. Tsinghua University Press.

[14] Theodoridis, S. (2020). Introduction. In *Machine learning: A Bayesian and optimization perspective* (2nd ed., pp. 1–17). Academic Press. https://doi.org/10.1016/B978-0-12-818803-3.00010-6

[15] Zhang, J., & Wang, S. (2024). Norms on fair use in training large models [in Chinese]. *ECUPL Journal, 27*(4), 20–33. https://doi.org/10.3969/j.issn.1008-4622.2024.04.002

[16] Chen, D. (2024). Legality analysis of artificial intelligence training other people's works under the background of data enclosure movement: From the perspective of technological process, innovation incentive and data equity [in Chinese]. *Electronics Intellectual Property, 2024*(10), 92–109. https://doi.org/10.3969/j.issn.1004-9517.2024.10.007

[17] Wan, Y. (2021). The dilemma and way out of the fair use system of copyright law in the era of artificial intelligence [in Chinese]. *Social Science Journal, 2021*(5), 93–102.

[18] Wang, Y. (2023). Copyright risks of text and data mining and their mitigation [in Chinese]. *Public Communication of Science & Technology, 15*(20), 137–139. https://doi.org/10.16607/j.cnki.1674-6708.2023.20.025

[19] Zhou, W., & Fei, Y. (2024). Fair use adjustment of the use of works by generative artificial intelligence creation [in Chinese]. *Science Technology and Law (Chinese-English Version), 2024*(3), 77–87. https://doi.org/10.19685/j.cnki.cn11-2922/n.2024.03.009

[20] Wu, H. (2020). *A study on the fair use system of copyright* (4th ed.) [in Chinese]. China Renmin University Press.

[21] Feng, X. (2006). A study of the incentive theory of copyright law: From the perspectives of economics, social welfare theory, and postmodernism [in Chinese]. *Science of Law (Journal of Northwest University of Political Science and Law), 2006*(6), 41–49. https://doi.org/10.16290/j.cnki.1674-5205.2006.06.006

[22] Xiong, Q., & Zhang, W. (2024). How copyright law regulate AI generated content: Global trends and local approaches [in Chinese]. *Intellectual Property, 2024*(11), 59–76.

[23] Leval, P. N. (1990). Toward a fair use standard. *Harvard Law Review, 103*(5), 1105–1136. https://doi.org/10.2307/1341457

[24] Yan, L. (2016). The rule of transformative use in US judicial practice and the implication [in Chinese]. *Intellectual Property, 2016*(6), 123–128. https://doi.org/10.3969/j.issn.1003-0476.2016.06.018

[25] Sag, M. (2012). Predicting fair use. *Ohio State Law Journal, 73*(1), 47–91. https://lawecommons.luc.edu/facpubs/180/

[26] Leaffer, M. A. (2010). *Understanding copyright law* (5th ed.). LexisNexis.

[27] Meng, Q., Li, X., & Yuan, D. (2019). The judging criteria for transformative use and its improvement path [in Chinese]. *Journal of Wuhan University of Technology (Social Sciences Edition), 32*(4), 115–123.

[28] Gordon, W. J. (1982). Fair use as market failure: A structural and economic analysis of the *Betamax* case and its predecessors. *Columbia Law Review, 82*(8), 1600–1657. https://doi.org/10.2307/1122296

[29] Xiong, Q. (2024). Adapting economic analysis of fair use to secondary use [in Chinese]. *Contemporary Law Review, 38*(1), 108–120. https://doi.org/10.3969/j.issn.1003-4781.2024.01.010

[30] Zhang, T. (2025). Copyright dilemmas in large-scale AI model training and adaptive regulatory pathways for legal harmonization [in Chinese]. *Modern Law Science, 47*(2), 189–208. https://doi.org/10.3969/j.issn.1001-2397.2025.02.13

[31] Zhang, X., & Bian, L. (2024). Copyright protection of generative artificial intelligence training corpus [in Chinese]. *Jingchu Law Review*, 2024(5), 77–90.

[32] Lin, P. K. (2023). Fair's fair: How public benefit considerations in the fair use doctrine can patch bias in artificial intelligence systems. *Indiana Journal of Law and Social Equality, 11*(2), 229–243. https://www.repository.law.indiana.edu/ijlse/vol11/iss2/1/

[33] Levendowski, A. (2018). How copyright law can fix artificial intelligence's implicit bias problem. *Washington Law Review, 93*(2), 579–630. https://digitalcommons.law.uw.edu/wlr/vol93/iss2/2