

A Study on the Normative Framework and Policy Dynamics of the Fair Use Doctrine under U.S. Copyright Law

Xiaojing Qin *

School of Law, Beijing Normal University, Beijing 100875, China;

* Corresponding Author. E-mail: xiaojing.qin@bnu.edu.cn

Abstract

Large models such as ChatGPT and Stable Diffusion rely heavily on automated scraping and machine learning of massive amounts of text, images, and code during pre-training. Whether this process, in its legal nature, constitutes an infringement of the copyright holder's exclusive rights or can be exempted through the fair use doctrine has become one of the most controversial frontier propositions in current intellectual property law. This article analyzes whether generative AI's scraping of training data constitutes fair use under U.S. copyright law. It systematically examines the application boundaries of the four-factor framework in AI scenarios, integrates the Copyright Office report and congressional legislative dynamics, reveals the tripartite game pattern among the judiciary, the administration, and the legislature, and provides methodological enlightenment for China's relevant institutional design.

Keywords

Artificial Intelligence, Fair Use, Copyright, China-U.S. Comparison, Generative AI.

1. Introduction

The rise of generative artificial intelligence is fundamentally reshaping the underlying logic of the copyright legal order. Large models such as ChatGPT and Stable Diffusion rely heavily on automated scraping and machine learning of massive amounts of text, images, and code during pre-training. Whether this process, in its legal nature, constitutes an infringement of the copyright holder's exclusive rights or can be exempted through the fair use doctrine has become one of the most controversial frontier propositions in current intellectual property law. [1] In terms of its institutional essence, fair use is a dynamic mechanism of interest balancing, whose function is to demarcate an elastic boundary between the private interest protection of copyright holders and the public interest of knowledge dissemination. However, when artificial intelligence "consumes" existing human intellectual achievements at an unprecedented scale and speed, whether the existing interpretive framework of fair use rules can continue to carry this balancing function warrants in-depth examination. Against this background, this article will systematically examine the core content of the U.S. fair use doctrine.

2. The Institutional Logic of the Four-Factor Framework and the Interpretive Tensions It Faces in the AI Era

The fair use doctrine is the most institutionally elastic normative device in the U.S. copyright law system. The system originated in nineteenth-century judicial precedents and was codified through Section 107 of the Copyright Act of 1976. Section 107 does not set closed constitutive requirements for fair use but requires courts to comprehensively weigh four statutory factors in individual cases, namely: the purpose and character of the use, the nature of the copyrighted work, the amount and substantiality of the portion used in relation to the work as a whole, and

the effect of the use upon the potential market for or value of the work, and to make an overall normative evaluation accordingly. The four factors are not isolated, checklist-style considerations but an organically linked evaluation system: the weight of any factor must be dynamically allocated within specific usage scenarios, and the factors often exhibit structural relationships of mutual traction and mutual restraint. It is precisely this openness that gives the system the ability to continuously adapt to technological iterations. From audio and video recording technology to the internet, and now to artificial intelligence, the core structure of the four-factor framework has remained substantially unchanged for half a century, and its institutional resilience itself constitutes an important object of observation for comparative law research. [2]

However, when generative AI incorporates copyrighted works into the training process at an unprecedented scale, the four-factor framework is facing the most severe interpretive tension in its institutional history: does large-scale, non-expressive data ingestion possess “transformative” purpose? [3] Does full copying in terms of quantity exceed the boundaries of fair use? [4] Does the indirect competitive effect of AI-generated, non-infringing content of the same type on the market for the original work constitute the “market harm” that the fourth factor is concerned with? These inquiries constitute the core problematic consciousness of the theoretical analysis and judicial practice examination that follow.

3. Theoretical Analysis of the Four Factors

3.1. Purpose and Character of the Use

This factor serves as the logical starting point of the four-factor analysis and bears the function of setting the tone for the entire fair use determination. [5] The Supreme Court systematically established the analytical paradigm for this factor in *Campbell v. Acuff-Rose Music, Inc.*, unfolding it into two interrelated review dimensions: transformative use and commercial nature. [6]

With respect to transformative use, Justice Souter in *Campbell* clearly pointed out that the core question is whether the accused use “merely supersedes the objects of the original creation” or “adds something new, with a further purpose or different character, altering the first with new expression, meaning, or message.” [6] The higher the degree of transformativeness, the more the weight of unfavorable factors such as commercial nature decreases in the overall weighing. The jurisprudential foundation of this inverse correlation lies in the fact that secondary uses that do not form market substitutes for the original work but instead open up expression spaces that the original work never touched are highly consistent with the constitutional purpose of copyright law to “promote the progress of science and useful arts.” [7, 8] *Google LLC v. Oracle America, Inc.* (hereinafter the “Google case”) further expanded the scope of this standard to the technical functional field, holding that Google’s copying of Java SE API declaring code was highly transformative because it served the functional purpose of “creating a new platform,” [9] indicating that the Supreme Court tends to define transformativeness from the perspective of functional purpose rather than formal appearance. However, *Andy Warhol Foundation v. Goldsmith* (hereinafter the “Warhol case”) imposed significant restrictions on this expansion trend: Justice Sotomayor emphasized that the determination of the first factor cannot rely solely on whether the secondary use “conveys a new meaning” but must examine whether “the specific purpose and character of the use” constitutes direct competition with the use of the original work. [10] From the expansion in *Campbell* to the contraction in *Warhol*, the evolution of the transformative use standard at the Supreme Court level has preset a tension-filled interpretive space for fair use analysis in AI training scenarios.

With respect to commercial nature, *Campbell* established a methodologically significant criterion: the key does not lie in whether the user’s sole motive is profit but in the broader

determination of the purpose and character of the use of the protected work. [6] The normative implication derived therefrom is that AI enterprises' behavior of scraping pirated sites or circumventing technological measures to obtain training data will bear a significant adverse presumption in the analysis of this factor.

3.2. Nature of the Copyrighted Work

This factor focuses on the attributes of the work itself being used. Its internal logic lies in the fact that the nature of the work determines the thickness of copyright protection, thereby indirectly affecting the strictness or leniency of fair use determination. *Harper & Row Publishers, Inc. v. Nation Enterprises* established two basic rules: the more fictional and expressive the work, the more its protection scope occupies the core territory of copyright law; the use of unpublished works constitutes a substantial weakening of the copyright holder's right of first publication. [11] However, as revealed by the *Google* case, this factor usually does not have decisive significance in the overall weighing, and its weight is clearly inferior to the first and fourth factors, [9] serving more as a moderating function.

3.3. Amount and Substantiality of the Portion Used

The *Google* case clearly pointed out that copying necessary or even complete portions to achieve a transformative purpose is acceptable; the key does not lie in absolute quantity but in "whether the user's legitimate purpose reasonably requires copying this amount." [9] *Harper & Row* pointed out that from the perspectives of quantity and substantiality, the degree of use should match the specific purpose of use, meaning that quantitative and qualitative analysis ultimately serves the judgment of purpose appropriateness. At the same time, even if it involves the "heart of the work," borrowing its information and ideas does not necessarily constitute copyright infringement, because copyright law protects expression rather than ideas themselves. [11] Projecting this doctrinal logic onto AI training scenarios, the controversial focus of this factor should shift from "whether the entire work was copied" to "whether full copying is necessary and proportionate for achieving the model's learning objectives"; otherwise, this factor will mechanically negate AI training behavior due to the technical characteristics of large model pre-training.

3.4. Effect on the Potential Market for or Value of the Work

Harper & Row established this factor as "the single most important element of fair use," [11] and its core function is to evaluate whether the accused use behavior constitutes a substantial adverse effect on the existing or potential market for the copyright holder's work. In AI training scenarios, the analytical complexity of this factor far exceeds traditional cases and can be unfolded from three progressive levels. First, direct substitution. If the AI model can reproduce the complete or a substantial portion of the original work, enabling the public to bypass the original work and directly obtain its core value, this constitutes the most powerful evidence for determining market harm. Second, licensing market harm. *Harper & Row* emphasized that fair use analysis must cover the impact on the "potential market," and the *Google* case further extended the scope of examination to potential markets that the copyright holder might license others to develop. [11, 9] With the advancement of commercial practices such as Shutterstock's signing of a six-year training data licensing agreement with OpenAI in 2023, a structural "AI training licensing market" has emerged, and AI enterprises' unauthorized training use will increasingly find it difficult to invoke "no licensing market exists" as a defense. Third, market dilution. If AI models can generate expressive content on a large scale that is similar in theme, style, or type to the original work and form indirect competition with it, such dilution effects also fall within the scope of market harm concerned by this factor. The judicial construction of this theory will be detailed in the analysis of the three major cases below.

4. Latest Policy Dynamics: The Copyright Office Report and Congressional Legislative Games

The interpretive space revealed by the above analysis has triggered fierce debate in U.S. policy-making and legislative games. The following provides a brief examination from two dimensions: the Copyright Office policy guidance and congressional legislative dynamics.

The report *Copyright and Artificial Intelligence—Part 3: Generative AI Training* published by the U.S. Copyright Office is the most comprehensive administrative policy guidance to date on the fair use issue of AI training. [12] The report adopts a cautious balanced position, holding on the one hand that AI training use of copyrighted works is likely to be transformative, while on the other hand explicitly rejecting the absolutist claim that AI training is essentially transformative. When AI output closely follows the creative intent of the input content, the Copyright Office tends to characterize such use as derivative rather than transformative. Regarding the fourth factor, the report systematically identifies three types of harm: direct sales loss, licensing market harm, and “market dilution,” and warns in a clear position: “The speed and scale at which AI systems generate content pose a serious dilution risk to the market for works of the same type.” At the institutional level, the report explicitly rejects supporting statutory compulsory licensing and instead advocates encouraging the market to develop creative licensing solutions on its own. The Copyright Office’s analysis of the application of the four factors, especially its limitation on the boundary of transformativeness of the first factor and its systematic exposition of the market dilution effect of the fourth factor, forms a dual echo of policy and adjudication with the judicial reasoning of the three major AI cases analyzed below. The focus of U.S. congressional legislation surrounding AI training data and copyright has shifted from information regulation centered on transparency disclosure in 2024, such as the *Generative AI Copyright Disclosure Act of 2024*, [13] to a direct challenge to the applicable boundaries of the fair use doctrine since 2025. The landmark event of this shift was the *AI Accountability and Personal Data Protection Act* jointly proposed by Senator Hawley (Republican) and Senator Blumenthal (Democrat) in July 2025. [14, 15] This bill creates a new federal tort cause of action, prohibiting AI enterprises from using copyrighted works for model training without “explicit, prior, informed, and meaningful consent,” and configures powerful remedies including compensatory damages, punitive damages, and injunctions. The industry has characterized it as “a direct attack on the fair use doctrine,” warning that if passed, it would substantially overturn recent rulings supporting fair use such as the *Bartz* and *Kadrey* cases. [16, 17] The *AI National Policy Framework* released by the White House in March 2026 stands in sharp contrast: the administration advocates that training behavior “does not violate copyright law” and recommends that Congress adopt a “non-intervention” stance, leaving specific application to the courts. [18]

5. Conclusion

The tripartite game pattern formed thereby among the judiciary, the administration, and the legislature provides two methodological enlightenments for China’s institutional reference: first, the “case-by-case adjudication” characteristic of the fair use doctrine faces a legitimacy challenge of insufficient normative supply when dealing with systematic industrial interest conflicts, which is an institutional cost that China needs to carefully assess when considering introducing similar systems; second, the game surrounding fair use is essentially an interest distribution conflict between the AI industry and content creators—choosing which institutional tool (judicial adjudication, administrative licensing, or legislative regulation) to configure this interest pattern itself constitutes a normative decision that needs to be made in conjunction with domestic industrial structure and cultural policy objectives.

References

- [1] Cooper, A. F., & Grimmelmann, J. (2025). The files are in the computer: On copyright, memorization, and generative AI. *Chicago-Kent Law Review*, 100, 141.
- [2] Lee, E. (2025). Fair use and the origin of AI training. *Houston Law Review*, 63, 105.
- [3] Sag, M. (2024). Fairness and fair use in generative AI. *Fordham Law Review*, 92, 1887.
- [4] Lucchi, N. (2024). ChatGPT: A case study on copyright challenges for generative artificial intelligence systems. *European Journal of Risk Regulation*, 15(3), 602–624.
- [5] Sag, M. (2023). Copyright safety for generative AI. *Houston Law Review*, 61, 295.
- [6] *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569 (1994).
- [7] U.S. Const. art. I, § 8, cl. 8.
- [8] Leval, P. N. (1990). Toward a fair use standard. *Harvard Law Review*, 103(5), 1105–1136.
- [9] *Google LLC v. Oracle America, Inc.*, 141 S. Ct. 1183 (2021).
- [10] *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508 (2023).
- [11] *Harper & Row, Publishers, Inc. v. Nation Enterprises*, 471 U.S. 539 (1985).
- [12] U.S. Copyright Office. (2025). Copyright and artificial intelligence, part 3: Generative AI training (Pre-publication version). U.S. Copyright Office.
- [13] Generative AI Copyright Disclosure Act of 2024, H.R. 7913, 118th Cong. (2024).
- [14] AI Accountability and Personal Data Protection Act, 119th Cong. (2025).
- [15] Office of Senator Josh Hawley. (2025, July 21). Press release on the AI Accountability and Personal Data Protection Act.
- [16] Chamber of Progress. (2025, July 21). Hawley-Blumenthal bill undermines innovation.
- [17] Miller Canfield PLC. (2025, August 4). Senators introduce legislation to curb use of personal data and copyrighted works for gen AI training.
- [18] The White House. (2026, March). AI national policy framework. The White House.