Journal of Sustainability, Policy, and Practice

Article

Rethinking Fair Use in the Age of AIGC: Balancing Copyright Protection and Technological Innovation

Qianwen Zhu [1,*]

[1] Faculty of Law, Minzu University of China, Beijing, China
* Correspondence: Qianwen Zhu, Faculty of Law, Minzu University of China, Beijing, China

Abstract: The rapid advancement of generative artificial intelligence (AI) is fundamentally reshaping the paradigm of content creation, yet the processes of data input, model training, and content generation entail significant risks of copyright infringement. Currently, Article 24 of China's Copyright Law, with its strict restrictions on eligible subjects and permissible purposes, proves increasingly ill-suited to the complex realities of artificial intelligence generated content (AIGC) creation. This article identifies three fundamental dilemmas within the existing legal framework: the explicit exclusion of non-natural persons from the scope of fair use, the inherent conflict between AIGC's technical requirements and the traditional "appropriate quotation" standard, and the inability of the teaching and scientific research provision to accommodate the massive scale of copying necessitated by machine learning algorithms. Drawing on a comprehensive comparative analysis of recent legislative approaches and judicial precedents in the United States and the European Union, this article proposes a strategic recalibration of the fair use doctrine. Specifically, it advocates for the introduction of conditional exception clauses tailored for AIGC, the systematic relaxation of subject and purpose restrictions, and the refined application of the international three-step test. Ultimately, these targeted legal adaptations seek to achieve a dynamic equilibrium between technological ethics and institutional rationality, ensuring both the robust protection of copyright holders' legitimate commercial interests and the reservation of sufficient institutional space for continued, sustainable innovation in the global artificial intelligence sector.

Keywords: generative ai; copyright; fair use; three-step test; data mining

1. Introduction

The rapid advancement of generative artificial intelligence (AI) has fundamentally reshaped the paradigm of content creation, accelerating the transition from professionally generated content and user-generated content to AI-generated content (AIGC). Generative AI refers to deep learning models that ingest input data and "learn" to produce outputs that mimic such data when duly prompted, a concept popularized by models such as GPT-3 and its successors. This technological leap, exemplified by systems like ChatGPT, Deepseek, and Ernie Bot, has been widely integrated into content production across textual, visual, and auditory domains, becoming an essential engine for innovation in the digital economy. However, the unprecedented capabilities of generative AI have precipitated profound challenges to existing copyright frameworks, particularly concerning the use of copyrighted materials for model training and the legal status of AI-generated outputs [1].

The copyright implications of AIGC manifest across three critical stages of the AI development lifecycle: data input, model training, and content generation. During the data curation and input phase, the collection and downloading of training datasets may involve reproducing substantial amounts of copyrighted works without authorization [2]. The model training stage presents even more complex issues, as algorithms analyze patterns, structures, and stylistic elements from protected materials to develop generative capabilities. Finally, the content output stage raises questions about whether AI-generated

Received: 09 May 2026
Revised: 15 June 2026
Accepted: 25 June 2026
Published: 28 June 2026

materials that bear similarity to training data may constitute infringing derivative works. These cumulative risks have sparked intense debate among scholars, practitioners, and policymakers regarding the applicability of copyright exceptions, particularly the fair use doctrine, to AIGC development and deployment.

The significance of this inquiry extends beyond doctrinal nicety to fundamental questions of innovation policy and economic development. As China's New Generation Artificial Intelligence Development Plan explicitly recognizes AI as a new engine of economic growth, the country has witnessed the emergence of numerous domestic models including Ernie Bot, Deepseek, and Doubao. These models depend on massive-scale data ingestion for training and optimization, creating an inherent tension between the closed authorization mechanisms traditionally favored by copyright law and the open data requirements of contemporary AI development. From a market regulation perspective, fair use serves as an institutional mechanism to correct market failures where transaction costs impede socially valuable uses of copyrighted works. In the AIGC context, the impracticality of obtaining individual licenses from millions of rightsholders for training data exemplifies precisely such market dysfunction.

Jurisdictional approaches to these challenges reveal significant divergence [3]. In the United States, the Copyright Office's recent report on copyright and artificial intelligence reaffirms that existing copyright law, particularly the fair use doctrine, can address AI training issues through case-by-case application of the four statutory factors, emphasizing transformativeness, commerciality, amount copied, and market effects. The report notes that while building training datasets using copyrighted works "clearly implicates the right of reproduction," the fair use defense may apply depending on the overall circumstances, with particular attention to whether outputs serve as functional substitutes for original works. The European Union has adopted a more structured approach through the Digital Single Market Directive, establishing text and data mining (TDM) exceptions calibrated to different use contexts, with one article permitting TDM for scientific research purposes by research organizations and another establishing a broader exception subject to rightsholder opt-out mechanisms. Meanwhile, China's judicial practice has begun to engage with these issues, as evidenced by the Beijing Internet Court's recognition that AIGC incorporating human intellectual input may qualify as protectable expression under copyright law. However, existing legislation, particularly Article 24 of China's Copyright Law, has yet to explicitly address whether AIGC development falls within the scope of fair use exceptions.

2. Theoretical Analysis of Fair Use in the AIGC Context

2.1. Defining AIGC and Identifying Copyright Risks

Generative artificial intelligence refers to deep learning models that ingest input data and "learn" to produce outputs that mimic such data when duly prompted. This definition, rooted in the technological essence of AI systems, captures the fundamental mechanism by which models such as GPT, Stable Diffusion, and their counterparts operate. The European Union employs the concept of TDM as a regulatory construct, emphasizing the balance between technological innovation and data rights, while U.S. discourse favors machine learning as a technology-neutral descriptor reflecting its pragmatic legal culture [4]. China's Interim Measures for the Management of Generative Artificial Intelligence Services defines generative AI as models and related technologies with content generation capabilities including text, images, audio, and video. For purposes of this article, AIGC refers specifically to content generated by such models and related technologies, distinguishing the output from the underlying technical processes.

The copyright implications of AIGC manifest across three critical stages of the AI development lifecycle. First, during the data curation and input phase, the collection and downloading of training datasets may involve reproducing substantial amounts of copyrighted works without authorization. Second, the model training stage presents more complex issues, as algorithms analyze patterns, structures, and stylistic elements from

protected materials to develop generative capabilities, potentially implicating the reproduction right [2]. Third, the content output stage raises questions about whether AI-generated materials that bear similarity to training data may constitute infringing derivative works, particularly where outputs regurgitate protected expression. These cumulative risks have generated intense debate regarding the applicability of copyright exceptions, particularly the fair use doctrine, to AIGC development and deployment.

2.2. The Concept and Evolution of Fair Use

Fair use constitutes a fundamental limitation on copyright, permitting unauthorized use of copyrighted works under specified conditions without liability. The doctrine originated in English common law, with a pivotal decision distinguishing between mere copying and use that adds new value through independent creation [5]. This principle was subsequently elaborated in American jurisprudence, most notably in an opinion that articulated three factors for evaluating fair use: the purpose and character of the use, the quantity and value of the materials used, and the effect on the original work's market. These factors were later codified in the U.S. Copyright Act, which added a fourth factor, the nature of the copyrighted work, and established the four-factor test that remains central to American fair use analysis.

Globally, two primary legislative approaches to fair use have emerged. The first, exemplified by the United States, adopts an open-ended, factor-based model that permits case-by-case adjudication based on statutory criteria. The second, followed by China and many civil law jurisdictions, employs a closed-list or rule-based approach enumerating specific exempted uses. China's Copyright Law, as amended in 2020, adopts this rule-based model in Article 24, which lists twelve specific circumstances constituting fair use, supplemented by a catch-all provision authorizing other circumstances provided by laws and administrative regulations. The article also incorporates the three-step test derived from the Berne Convention, requiring that any limitation on exclusive rights be confined to certain special cases, not conflict with a normal exploitation of the work, and not unreasonably prejudice the legitimate interests of the rightsholder [6].

2.3. Justifying Fair Use Application to AIGC

Scholarly discourse on whether AIGC should benefit from fair use protection divides into affirmative and negative positions. Proponents argue that extending fair use to AIGC promotes socially beneficial outcomes including text and data mining, information dissemination, and technological innovation. They emphasize that the transformative nature of AI training, extracting patterns and statistical relationships rather than reproducing expressive content, aligns with the core justifications for fair use [7]. Skeptics counter that excessive accommodation of AIGC risks technological abuse and may discourage human creativity by diminishing the economic incentives that copyright is designed to protect.

From an industrial development perspective, fair use aligns strategically with national artificial intelligence objectives. China's New Generation Artificial Intelligence Development Plan explicitly recognizes AI as a new engine of economic growth and emphasizes deep learning, cross-boundary fusion, and human-AI synergy as characteristic features of next-generation AI. Domestic models including Ernie Bot, Deepseek, and Doubao depend on massive-scale data ingestion for training and optimization. Requiring individual licenses from millions of rightsholders for training data would impose transaction costs so prohibitive as to impede technological development entirely [8]. As the 14th Five-Year Plan identifies AI as a strategic emerging industry, a closed authorization mechanism ill-suited to generative AI's data requirements would undermine industrial competitiveness.

From a market regulation perspective, fair use serves as an institutional mechanism to correct market failures where transaction costs impede socially valuable uses. Professor Wendy J. Gordon's seminal analysis identifies transaction costs as critical barriers to efficient resource allocation: where such costs exceed expected benefits, mutually beneficial exchanges cannot occur through normal market mechanisms. Copyrighted

works exhibit public goods characteristics, making ex ante prevention of unauthorized use difficult and rendering consensual negotiations often infeasible. Fair use reduces these transaction costs by permitting socially valuable uses without requiring formal licensing, thereby promoting dissemination, encouraging follow-on creation, and facilitating optimal resource allocation [9]. In the AIGC context, the impracticality of securing millions of individual licenses exemplifies precisely such market dysfunction.

3. The Tripartite Predicament of AIGC under Existing Fair Use Provisions

3.1. Subject Predicament

Article 24(1)(1) of China's Copyright Law exempts from liability the use of published works "for the purposes of personal study, research, or appreciation." This provision, which constitutes the core of China's "personal use" exemption, presents insurmountable obstacles to AIGC applications. Textually, the limitation to "personal" use confines the exemption to natural persons exercising independent cognitive agency. Legal persons, as artificial constructs lacking the capacity for study or research in the human sense, fall outside its ambit. The "personal" requirement cannot be interpreted to encompass legal persons or AI systems themselves, as these lack the autonomous cognitive capacity necessary for activities such as learning and research that are inherently human endeavors.

This textual limitation assumes practical significance given the identity of AIGC developers [10]. The principal actors in generative AI development, technology enterprises such as Huawei, Baidu, and Alibaba are unmistakably legal persons whose research activities are inseparable from commercial objectives. The 2025 Suzhou Intermediate People's Court decision in an AI-generated image copyright dispute illustrates this tension: the court rejected copyright claims where users merely provided initial prompts without demonstrating substantial creative contribution through parameter adjustment and model optimization. While this case addressed copyrightability rather than fair use, it underscores the courts' reluctance to extend legal protections to purely machine-driven processes absent meaningful human intervention.

The practical consequence is that commercial AI developers cannot invoke the personal use exemption. An interpretive expansion of "personal" to encompass legal persons would not only strain statutory language beyond its breaking point but also create perverse incentives for commercial entities to cloak systematic copyright infringement in the garb of technological research. As the Shenzhen Nanshan District Court recognized in the 2019 Tencent AI-generated financial report case, legal persons may qualify as "fictitious authors" where algorithm selection and data analysis reflect creative labor under corporate direction, but this recognition of legal person authorship for AI-generated outputs does not automatically extend to legal person eligibility for personal use exceptions during the training phase [11].

3.2. Conduct Predicament

Article 24(1)(2) of China's Copyright Law permits the "appropriate quotation" of published works for purposes of "introduction, commentary, or explanation." While this provision avoids the subject-matter difficulties of the personal use exemption, it imposes purposive and quantitative constraints fundamentally at odds with the operational logic of AI-generated content (AIGC).

The purposive limitation proves particularly problematic [12]. News organizations and libraries increasingly employ AI for content generation, with examples such as the use of AI-generated video content by certain platforms. Such uses, while serving legitimate public information functions, simultaneously generate commercial value through the creation of paid knowledge products and cross-domain data exploitation. This dual-purpose character, simultaneously public-serving and commercially valuable, resists easy categorization under the "introduction, commentary, or explanation" rubric. Recent guidance emphasizes that fair use determinations require careful analysis of whether AI outputs serve as functional substitutes for original works, with transformative purpose weighing heavily in the analysis.

Quantitative constraints present equally intractable difficulties. The "appropriate" limitation presupposes bounded quotation, yet AIGC's underlying logic demands maximal data ingestion to optimize model performance. The technical imperative of more data leading to better models collides directly with legal expectations of measured taking. Moreover, the opacity of algorithmic processing, often referred to as the black box problem, renders it practically impossible for users to ascertain whether outputs incorporate protected expression from training materials [13]. Legal disputes have demonstrated that plaintiffs face substantial evidentiary hurdles in proving that specific copyrighted works were used in training and contributed to allegedly infringing outputs.

The fundamental incompatibility lies in the distinction between traditional quotation, which uses portions of a work to illustrate or comment upon that work, and AI training, which uses entire works to extract patterns and statistical relationships [3]. This distinction reflects the difference between expressive and non-expressive uses, with the latter falling outside the core concerns that copyright law seeks to address.

3.3. Quantitative Predicament

Article 24(1)(6) of China's Copyright Law exempts "the translation, adaptation, compilation, broadcasting, or reproduction in small quantities of published works for use by teaching or scientific research personnel." This provision, like its counterparts, imposes restrictions that artificial intelligence-generated content (AIGC) cannot satisfy [5].

The quantitative limitation to "small quantities" proves fatal. Model training requires full-text reproduction and analysis of works at scales that overwhelm any plausible interpretation of "small." The text and data mining (TDM) exceptions in European Union law were designed for automated analytical extraction for scientific or informational purposes, not the large-scale reproduction, internalization, and expressive re-use of works characteristic of generative AI training. Article 3 of the Digital Single Market (DSM) Directive is limited to non-commercial research by research organizations and cultural heritage institutions, while Article 4's opt-out mechanism, based on non-standardized signals, exacerbates uncertainty without ensuring transparency or fair compensation [1].

Subject-matter and purposive restrictions compound the difficulty [4]. The Chinese exemption applies only to teaching and research activities conducted by educational and scientific research institutions serving public, non-commercial purposes. Yet AIGC development is dominated by commercial enterprises whose activities, even when employing scientific personnel, constitute employment-related acts within the meaning of Article 170 of the Civil Code, fundamentally different from the non-profit, public-interest research contemplated by the exemption.

The European approach illustrates the limitations of structured exceptions for AI training. Article 3 of the DSM Directive permits TDM for scientific research but restricts eligibility to research organizations and cultural heritage institutions, excluding commercial actors entirely. Article 4 establishes a broader exception for TDM purposes but subjects it to rightsholder opt-out rights, creating significant uncertainty for AI developers regarding the legality of their training datasets. The pending reference in Like Company v. Google Ireland seeks clarification on whether large language model (LLM) training constitutes reproduction within the meaning of Article 2 of the InfoSoc Directive and whether the TDM exception applies to commercial AI training.

The inadequacy of existing quantitative limitations reflects a deeper conceptual mismatch. Traditional fair use analysis examines the amount and substantiality of the portion used relative to the copyrighted work as a whole, assuming bounded taking from identifiable works. AI training inverts this paradigm: it requires taking from millions of works to extract patterns that transcend any individual work's expressive content. This "non-expressive" use challenges the very framework on which quantitative limitations rest [8].

4. Comparative Perspectives

4.1. The United States

United States copyright law adopts an open-ended, factor-based approach to fair use, codified in 17 U.S.C. § 107. The four statutory factors—purpose and character of use, nature of the copyrighted work, amount and substantiality of the portion used, and effect on potential markets—provide flexible criteria for case-by-case adjudication [13]. This flexibility has proven particularly valuable in addressing novel technological applications, with the transformative use doctrine emerging as the dominant analytical framework following its articulation by Judge Leval and endorsement by the Supreme Court in Campbell v. Acuff-Rose Music.

The landmark Authors Guild v. Google decision exemplifies this approach. The Second Circuit upheld Google's mass digitization of copyrighted books, finding that the creation of a searchable database and provision of snippet views served purposes fundamentally different from the original expressive purposes of the scanned books, thereby constituting transformative use. The court emphasized that the use "augments public knowledge by making available information about Plaintiffs' books without providing the public with a substantial substitute for" the originals. This decision established two principles of particular relevance to AIGC: first, that mass reproduction may qualify as fair use under appropriate circumstances; and second, that commercial purpose, while relevant, does not preclude a finding of fairness where use serves socially valuable functions distinct from the original work's purpose.

Recent developments in U.S. litigation have extended this analysis directly to generative AI training. In June 2025, the Northern District of California issued the first fair use decisions in GenAI litigation, both finding that training on copyrighted books constituted fair use. In Bartz et al. v. Anthropic, the court granted summary judgment for Anthropic, holding that using books to train the Claude LLM was exceedingly transformative [8]. The court reasoned that Anthropic used the works to iteratively map statistical relationships between text fragments so that Claude could return new text outputs as if it were a human reading prompts and writing responses, rather than providing users with excerpts from the books themselves. Drawing an analogy to Google Books, the court emphasized that guardrails ensuring no infringing content reached users supported the transformative purpose finding.

In Kadrey et al. v. Meta, another California federal court reached the same conclusion regarding Meta's LLaMA model, albeit through different reasoning. While acknowledging that the case involved the creation of a product with the ability to severely harm the market for the works being copied, the court found that plaintiffs presented no evidence of actual market dilution, causing the fourth factor to weigh in Meta's favor. These decisions demonstrate that U.S. courts may focus on different fair use factors—transformativeness in Bartz, market harm in Kadrey—while reaching consistent outcomes. The U.S. Copyright Office's Part III Report on Copyright and Artificial Intelligence, released in May 2025, confirms this approach, concluding that some uses of copyrighted works for generative AI training will qualify as fair use, and some will not, with analysis required on a case-by-case basis.

4.2. The European Union

The European Union has pursued a more structured approach through the Digital Single Market Directive 2019/790. Articles 3 and 4 establish TDM exceptions calibrated to different use contexts. Article 3 permits TDM for scientific research purposes by research organizations and cultural heritage institutions, without the possibility of contractual override, but excludes commercial actors entirely. Article 4 establishes a broader exception for TDM purposes, applicable to any user, but subject to an opt-out mechanism permitting rightsholders to reserve their rights in an appropriate manner, such as machine-readable means. The AI Act 2024/1689 reinforces these provisions, requiring providers of general-purpose AI models to implement policies respecting copyright and to make publicly available "sufficiently detailed summaries" of training data.

The EU framework's operation is subject to the three-step test incorporated through Article 7(2) of the DSM Directive, which refers to Article 5(5) of the Information Society

Directive. This requires that exceptions apply only "in certain special cases which do not conflict with a normal exploitation of the work or other subject-matter and do not unreasonably prejudice the legitimate interests of the rightholder." Scholarly analysis suggests that arguments for finding a conflict with normal exploitation based on market dilution by AI outputs may be premature. As one expert observes, where AI output reproduces only ideas, concepts, and styles, matters excluded from copyright protection under the idea or expression dichotomy, it is "impossible to establish a conflict with 'a normal exploitation of the work'" because rightsholders "do not have the right to prohibit competition at idea, concept and style level."

The pending reference in Like Company v. Google Ireland (C-250/25) promises to clarify the scope of EU exceptions as applied to generative AI. The Budapest District Court referred four questions to the CJEU in April 2025, seeking guidance on whether chatbot outputs reproducing press articles constitute communication to the public, whether LLM training constitutes reproduction under EU law, if such reproduction falls within the Article 4 TDM exception, and liability for user-prompted outputs [3]. The referring court explicitly asks whether the TDM exception, designed for analytical extraction of patterns, can legitimately extend to generative AI training that synthesizes new expressions. A ruling favoring rightsholders could require AI developers to obtain licenses for both training and deployment, while a ruling favoring AI developers may confirm that generative outputs are sufficiently transformative to fall outside reproduction and communication rules.

The practical implementation of EU rules faces significant challenges. The opt-out mechanism's effectiveness depends on machine-readable reservations, yet no single, uniformly adopted technical standard exists. The Robot Exclusion Protocol (robots.txt) and TDM Reservation Protocol have serious limitations, and the European Commission has ordered a feasibility study for a central registry of TDM opt-outs. Member state implementation varies, with some expressly requiring machine-readable means for online content while others do not. These uncertainties, combined with the absence of statutory licensing schemes in any examined jurisdiction, mean that compensation for rightsholders currently depends entirely on contractual arrangements based on property rights [4].

4.3. Lessons for China

Comparative analysis yields three insights relevant to Chinese law reform. First, jurisdictions with advanced AI industries have generally adopted permissive stances toward TDM for AI training, whether through flexible fair use adjudication in the United States or structured exceptions in the European Union [9, 10]. The U.S. approach prioritizes flexibility and case-by-case assessment, while the EU emphasizes legal certainty through clearly defined exceptions subject to opt-out mechanisms.

Second, permissiveness is typically conditioned on safeguards designed to protect rightsholders' legitimate interests [3]. The U.S. emphasis on transformative purpose and market effects ensures that use qualifying as fair use serves purposes distinct from the original works. The EU's opt-out mechanism and transparency requirements provide rightsholders with meaningful opportunities to protect their works from exploitative uses. As recognized in judicial decisions, even where training qualifies as fair use, maintaining unauthorized copies for undefined purposes may not.

Third, the trend favors technologically neutral approaches that accommodate innovation without requiring constant legislative revision. The U.S. fair use doctrine's adaptability to novel technologies, demonstrated from Google Books to Anthropic, offers a model for judicial evolution. The EU's structured exceptions, while providing greater predictability, face challenges in accommodating technologies not contemplated at the time of drafting, precisely the issue raised in cases regarding whether TDM exceptions designed for analytical purposes can properly extend to generative AI.

These insights suggest that China need not choose between technological development and copyright protection [9]. Rather, the challenge lies in designing mechanisms that permit socially valuable AI training while providing rightsholders with

meaningful opportunities to protect their works from exploitative uses. The absence of any TDM exception in China's current Copyright Law places China at the restrictive end of the international spectrum, alongside the United Kingdom following its abandonment of proposed TDM expansion. However, judicial recognition that AI training may qualify as fair use under appropriate circumstances demonstrates that explicit exceptions are not always necessary.

5. Toward a Calibrated Fair Use Framework for AIGC

5.1. Legislative Refinement

The most direct legislative response to the challenges identified above would introduce a specialized exception for AIGC as a thirteenth enumerated case under Article 24 of China's Copyright Law. Rather than attempting to stretch existing provisions beyond their textual and conceptual limits, a targeted exception calibrated to the distinctive characteristics of AI training would provide legal certainty while accommodating technological innovation [9]. Japan's approach in its Copyright Act offers a useful model for such legislation.

The Japanese provision, introduced in 2018 and effective from 2019, comprehensively permits exploitation of a work that is aimed at neither enjoying nor causing another person to enjoy the work. This non-enjoyment purpose exception covers any exploitation for TDM, including for commercial purposes, without restriction as to means or scale. The underlying theory relates to the nature of copyright itself: exploitation not for enjoyment purposes falls outside the inherent scope of copyright because it does not prejudice the opportunities of copyright holders to receive compensation. This approach differs fundamentally from both the U.S. fair use doctrine and the U.K. fair dealing provisions, offering a more targeted solution for uses, such as AI training, that do not implicate the core expressive purposes that copyright protects [11].

Adapting this model to the Chinese context, a provision might permit the reproduction and analysis of works for the purpose of generating new knowledge or information through data processing, including but not limited to model training, data mining, and information analysis. The provision should be subject to conditions requiring that the use serves public interest purposes, does not unreasonably prejudice the legitimate interests of rightsholders, and complies with any reservation of rights communicated through appropriate technical means. This formulation captures the essential character of AI training, the extraction of patterns and statistical relationships from works rather than the reproduction of expressive content for communicative purposes, while building in protections for rightsholders through the opt-out mechanism.

The public interest purposes condition recognizes that not all AI training serves socially valuable ends. While TDM provisions permitting use of human works without financial compensation raise questions of parasitic use, the payment of equitable remuneration changes the equation. If authors and rightsholders are remunerated appropriately for the use of their works in AI training, the process can no longer be described as parasitic and unduly harmful. This insight points toward the possibility of combining TDM freedom with an obligation to pay equitable remuneration through statutory licensing or levy systems, an approach explored further below [10].

The opt-out mechanism merits particular attention [9]. Following the EU model in Article 4(3) of the DSM Directive, rightsholders should be able to reserve their rights through appropriate means, such as machine-readable protocols. However, the practical implementation of opt-out systems faces significant challenges. The Robot Exclusion Protocol and TDM Reservation Protocol have serious limitations, and no single, uniformly adopted technical standard yet exists. The European Commission has ordered a feasibility study for a central registry of TDM opt-outs, recognizing that fragmented, non-standardized reservation mechanisms exacerbate uncertainty without ensuring transparency or fair compensation.

China's National Technical Committee 260 on Cybersecurity has already provided guidance in this area through the "Basic Security Requirements for Generative AI Service," which recommends respecting the Robots Exclusion Protocol and other technical restrictions. Building on these standards, legislation should authorize the Cyberspace Administration of China, in consultation with the National Copyright Administration, to designate approved technical means for rights reservations and maintain a public registry of such means to ensure accessibility and interoperability [2].

5.2. Subject Expansion

Legislative reform should expand the scope of eligible subjects to encompass entities engaged in bona fide AIGC development, regardless of their commercial or non-commercial character. The diverse participation landscape observed in Chapter 3, encompassing technology giants, small and medium-sized enterprises, research institutions, and individual developers, reflects technological realities that legal rules must accommodate [9]. A subject-matter restriction limited to traditional research institutions would exclude precisely those actors best positioned to advance AI technology and translate research into practical applications.

On purpose, legislation should abandon the non-commercial prerequisite. As demonstrated in other jurisdictions, commercial purpose does not preclude fair use where use serves socially valuable functions distinct from the original work's purpose. Both the Google Books and Anthropic cases involved commercial entities whose AI training and digitization projects served research and discovery functions fundamentally different from the expressive purposes of the works used. The relevant inquiry concerns the nature and consequences of use, not the user's profit motivation.

The 2024 Hangzhou Internet Court decision in a case involving an AI product trained on comics offers suggestive guidance. The court suggested that using works for training without directly harming the copyright owner's market may be seen as permissible, aligning with a more liberal approach emphasizing transformative nature and lack of direct market competition [6]. While this decision does not establish binding precedent, it signals judicial openness to considering the distinctive characteristics of AI training in fair use analysis.

However, relaxing subject and purpose restrictions must be accompanied by appropriate safeguards. The Indian Department for Promotion of Industry and Internal Trade's December 2025 working paper on AI and copyright offers an instructive model. Rejecting both blanket exemptions and pure licensing approaches, the Committee proposed a hybrid model under which: AI developers receive a blanket licence for use of all lawfully accessed content for training purposes without individual negotiations; royalties become payable only upon commercialisation of AI tools, with rates set by a government-appointed committee; and a centralised mechanism handles royalty collection and distribution to reduce transaction costs and provide legal certainty. This approach recognizes that the timing and mechanism of compensation matter as much as its existence, deferring payment obligations until commercialisation reduces barriers to entry while ensuring rightsholders share in the value created by AI systems trained on their works.

For China, a similar approach could be calibrated to domestic institutional frameworks. China's existing copyright collective management organizations for music, text, and images could be leveraged to broker deals allowing AI firms to use large catalogues of works while paying fair remuneration to rightsholders. The transaction costs of individual licensing at the training stage would be prohibitive, but collective licensing mechanisms can aggregate rights and distribute royalties efficiently. The key design challenge lies in ensuring that such systems do not disadvantage small developers or individual authors, concerns that can be addressed through graduated fee structures and mandatory pass-through requirements ensuring that licensing revenues reach actual creators rather than being captured entirely by intermediaries.

5.3. Standard Clarification

Application of the three-step test, specific case, no conflict with normal exploitation, no unreasonable prejudice to legitimate interests requires adaptation to AIGC's distinctive characteristics. The first step presents the least difficulty. AI training for generative purposes should qualify as a specific case given its technological necessity and social utility in promoting innovation. The certain special cases requirement does not demand narrow, fact-specific exceptions; it permits appropriately defined categories of use that serve legitimate purposes. The proposed conditional exception clause satisfies this requirement by specifying the types of use permitted and the conditions attached.

The second step, no conflict with normal exploitation, requires more careful analysis. Conflict arises where use substantially substitutes for original works in their markets. AI training that internalizes works for pattern extraction, without making works publicly accessible, does not substitute for original works. However, where outputs reproduce protected expression or serve as market substitutes, as alleged in certain legal disputes, conflict may arise. The inquiry requires attention to functional substitution, not merely formal copying.

Professor Senftleben's analysis of the normal exploitation test offers crucial guidance. In cases where AI output only reproduces ideas, concepts, and styles, it is impossible to establish a conflict with "a normal exploitation of the work" because rightsholders do not have the right to prohibit competition at the idea, concept, and style level under the fundamental idea or expression dichotomy recognized in international agreements. AI productions that offer style imitations as alternatives to human works may have a diffuse disruptive effect on the market for human literary and artistic productions, but this unspecific competition at the abstract level of ideas, concepts, and styles cannot be deemed sufficient to assume a conflict with normal exploitation of a specific work. Unless an author's individual creative choices have been copied, rightsholders cannot demonstrate that the market for a specific work has been usurped.

The third step, no unreasonable prejudice, offers the most promising avenue for protecting rightsholders' legitimate interests while accommodating AI development. From the adoption of the three-step test at the 1967 Stockholm Conference for the Revision of the Berne Convention, it has been recognized that equitable remuneration can serve as a tool to reduce unreasonable prejudice to a permissible reasonable level. In the copyright or AI debate, the payment of remuneration is a central balancing instrument [2, 3].

For lawmakers seeking to reconcile AI innovation with copyright protection, the unreasonable prejudice test offers additional breathing space. While a TDM provision permitting use of human works in AI training without any financial compensation raises questions of parasitic use and may cause unreasonable prejudice, the payment of equitable remuneration changes the equation. If authors and rightsholders are remunerated appropriately for the use of their works in AI training, the TDM process can no longer be described as parasitic and unduly harmful.

This insight opens the door to statutory licensing or levy solutions [5]. Such systems can be implemented at the AI training stage or, alternatively, the obligation to pay fair remuneration can be imposed when fully developed AI systems are finally brought to the market. Regardless of the reference point, AI training remuneration can be channelled to collecting societies which ensure with appropriate repartitioning schemes that collected money is distributed fairly between authors and rightsholders.

From the perspective of authors fearing displacement effects, an inescapable obligation to pay equitable remuneration has the welcome side effect that it makes the use of GenAI systems more expensive. Unless GenAI providers are willing to pay the remuneration out of their own pocket, they must pass on costs to users of their systems, reducing the advantage of lower production costs and enhancing the chances of human authors to compete. AI remuneration can be set at a level that counterbalances cost savings, transforming AI content revenue into human content revenue.

In practical implementation, the assessment should move beyond purely quantitative analysis. Drawing on the idea or expression dichotomy, use of works for non-expressive learning, extracting unprotected ideas, facts, and patterns should not constitute

unreasonable prejudice. Use for expressive learning, reproducing protected expression in outputs may do so. The analysis must also consider opt-out mechanisms and transparency requirements as means of minimizing prejudice while preserving use opportunities. The "Basic Security Requirements for Generative AI Service" (GB/T 45654-2025) and "Security Specification for Generative AI Pre-Training and Fine-Tuning Data" (GB/T 45652-2025) already establish baseline expectations for data governance that can inform this analysis [2].

6. Conclusion

This study has examined the applicability of fair use to AIGC under China's Copyright Law. The analysis reveals a tripartite predicament: the "personal use" exemption excludes legal persons from eligibility; the "appropriate quotation" standard imposes purposive and quantitative constraints at odds with AI training requirements; and the "teaching and research" exception cannot accommodate machine learning's scale of copying.

Comparative analysis yields three insights: leading jurisdictions adopt permissive stances toward TDM; permissiveness is conditioned on safeguards protecting rightsholders; and technologically neutral approaches accommodate innovation without constant legislative revision.

Building on these insights, this article proposes a three-element framework: conditional exception clauses modeled on Japan's "non-enjoyment" purpose exception; expanded eligibility for all bona fide developers with collective licensing safeguards; and refined three-step test application recognizing that non-expressive pattern extraction does not conflict with normal exploitation when equitable remuneration is provided.

Deferring compensation until commercialization reduces barriers to entry while ensuring rightsholders share in value created. Equitable remuneration transforms potential prejudice into permissible accommodation and, by increasing AI system costs, enhances human creators' competitive position.

Ultimately, by designing mechanisms that permit socially valuable AI training while ensuring rightsholders share in value created, China can achieve the dynamic equilibrium between technological innovation and copyright protection that the digital age demands.

References

1. P. Samuelson, "Fair use defenses in disruptive technology cases," 2023.
2. C. L. Saw and B. Z. Y. Tan, "Unpacking copyright infringement issues in the GenAI development lifecycle and a peek into the future," *Computer Law & Security Review*, vol. 58, p. 106163, 2025.
3. E. Lee, "Fair use and the origin of AI training," *Houston Law Review* (forthcoming 2025), Santa Clara Univ. Legal Studies Research Paper, no. 5253011, 2025.
4. W. Jon, "Prompting creativity: Tiered approach to copyright protection for AI-generated content in the digital age," *Media and Communication*, vol. 13, 2025.
5. W. J. Gordon, "Fair use as market failure: a structural and economic analysis of the 'Betamax' case and its predecessors," *Columbia Law Review*, vol. 82, no. 8, pp. 1600-1657, 1982.
6. P. Goldstein and P. B. Hugenholtz, *International copyright: principles, law, and practice*, Oxford University Press, 2019.
7. T. Chen, "Language Models' Verbatim Copying: Copyright Infringement Analysis through the Lens of The New York Times Co. v. Microsoft Corp., OpenAI, Inc. et al.," *Microsoft Corp., OpenAI, Inc. et al.*, Apr. 3, 2025, vol. 43.
8. B. Sobel, "Artificial intelligence's fair use crisis," 2017.
9. D. Nimmer, "Copyright in the Dead Sea scrolls: authorship and originality," *Houston Law Review*, vol. 38, no. 1, 2001.
10. J. C. Charlesworth, "Generative AI's illusory case for fair use," *Vanderbilt Journal of Entertainment & Technology Law*, vol. 27, no. 3, p. 323, 2025.
11. V. Liu, "Kadrey v. Meta Platforms, Inc. No. 3:23-cv-03417-Vc (ND Cal. June 25, 2025)," *Intell. Prop. & Tech. LJ*, vol. 30, p. 78, 2025.
12. J. P. Quintais, "Generative AI, copyright and the AI Act," Available at SSRN 5021148, 2024.
13. M. Senftleben, "TDM, GenAI and the copyright three-step test," *International Review of Intellectual Property and Competition Law (IIC)*, vol. 57, pp. 67-107, 2026.