Proceeding of ICGPSH 2026 Symposium: EdTech & AI in Learning: Large Language Models in Business, Politics, and Humanities
DOI: 10.54254/2753-7048/2026.LD33864

Ethical Embedding and Institutional Reconstruction of Intellectual Property Governance for Generative Artificial Intelligence

Kangyi Wang[1*], Renfu Liu[1], Ruisheng Wang[1], Bo Pang[1], Tianlang Zhou[2]

[1]*School of Law, Southwest University of Science and Technology, Mianyang, China*
[2]*School of River and Ocean Engineering, Chongqing Jiaotong University, Chongqing, China*
**Corresponding Author. Email: wang-kangyi@outlook.com*

Abstract. Generative artificial intelligence poses systematic challenges to the traditional human-centrism-based intellectual property system in two dimensions: the fair use of training data and the copyrightability of generated content. It exposes the deep ethical tension between technological innovation and the protection of creators' rights and interests. Adopting a functionalist comparative law approach, this paper systematically analyzes the differentiated regulatory paths regarding the use of AI training data and the copyrightability of generated content in the United States, the United Kingdom, Japan and China. Combined with utilitarian incentive theory and personality right ethics, it explores the ethical limitations of paradigms such as the Absolute Control Doctrine and the Substantial Contribution Doctrine in terms of transparency, fairness and liability attribution. By establishing the theoretical benchmark of "creativity ecosystem integrity", this paper constructs a hierarchical consent mechanism, an ethical review system for copyrightability, as well as an ethical insurance and risk-sharing mechanism. It aims to achieve a dynamic balance between ensuring fair participation of human creators in value distribution and promoting technological innovation, and form a normative framework for global artificial intelligence intellectual property governance.

Keywords: Generative AI, Training Data Ethics, Comparative Intellectual Property Law, Human-Centrism

1. Systematic challenges of generative AI to the intellectual property system

John Locke laid the ethical foundation of modern property rights in *Two Treatises of Government*: "Every man has a property in his own person; this nobody has any right to but himself. The labour of his body and the work of his hands, we may say, are properly his." [1] Revolutionary technological breakthroughs are often accompanied by adaptive adjustments to production relations. As a core driving force of the new round of scientific and technological revolution, generative artificial intelligence is reconstructing the technical foundation of knowledge production through the in-depth integration of computing power, algorithms and data, giving rise to a new type of production relation of "human-machine collaboration". At present, the "nonexpressive use" of

DOI: 10.54254/2753-7048/2026.LD33864

massive works by generative AI has triggered deep conflicts with the traditional intellectual property system.

Immanuel Kant once stated: "Act in such a way that you treat humanity, whether in your own person or in the person of any other, never merely as a means to an end, but always at the same time as an end." [2] The existing intellectual property system is built on a fundamental philosophical premise: the term "intellectual" refers to human intelligence, and the subjects of creation and invention are presumed to be natural persons with independent consciousness and moral responsibility [3]. This human-centric tradition manifests itself in copyright law as the requirement for "originality": creators must inject personal thoughts, emotions and judgments into their works. However, as AI systems increasingly demonstrate the capacity for creative selection based on opaque algorithms, the exclusive attribution of intelligence to humans has shifted from a self-evident truth to a proposition awaiting proof [4].

This raises two key sub-questions: how to determine whether data mining exceptions in the training stage constitute "ethical exploitation" of creators, and what ethical standards to adopt to recognize the copyright of generated content.

2. Theoretical tension between ethics and intellectual property

The legitimacy of modern copyright law is mainly justified by two traditional theories: Locke's labor theory of property rights provides a natural law foundation, while the utilitarian tradition offers an efficiency-oriented justification.

Locke's labor theory of property rights holds that individuals acquire moral property rights over the fruits of their labor by investing labor in natural resources. Extending this to the field of intellectual creation, authors who produce works through mental labor deserve copyright protection. From a consequentialist perspective, utilitarianism argues that the legitimacy of the copyright system lies in its ability to maximize social welfare. Providing economic incentives to creators is a means, and promoting knowledge production and dissemination is the end [5].

Unlike the Anglo-American legal system, which emphasizes economic incentives, author's rights in the civil law system are rooted in personality right ethics. It advocates that a work is an externalization of the author's personality, and the primary purpose of copyright protection is not economic incentives, but the protection of creators' spiritual and personality interests. For example, Article 2 of the *Japanese Copyright Act* defines a "work" as a "creative expression of thoughts or emotions", explicitly requiring that a work must reflect human thoughts or emotions [6]. In the causal chain of copyright law, human causal contribution is a prerequisite for granting copyright protection to a work.

3. Ethical and legal comparison of training data usage

3.1. Spectrum analysis of global regulatory approaches

Generative AI's reliance on massive training data creates a normative paradox at the input end. There is a structural irreconcilability between the scale of data acquisition required for technological innovation and the individual authorization logic presupposed by copyright law. This is not merely a legal technical issue, but one that spans ethics, law and political economy. Traditional copyright rules such as licensed use, fair use and statutory licenses have all failed to varying degrees when faced with the massive data demand of generative AI [7]. However, there is a lack of systematic

DOI: 10.54254/2753-7048/2026.LD33864

theoretical integration in the comparative analysis of how different jurisdictions respond to this dilemma based on different value presuppositions.

3.1.1. Lax model: Japan's "non-appreciation purpose" exception

Japan's institutional response centers on Articles 30-4 and 47-5 of its *Copyright Act*, forming one of the most lenient legal frameworks globally for the use of AI training data [8]. The Japanese Copyright Act establishes "non-appreciation purpose" as the core criterion for determining fair use. Its premise is that the reproduction of works in machine learning is not for the purpose of "appreciating the expression of the work", but for extracting statistical laws and information patterns, and thus does not cause substantial harm to the interests protected by copyright [9]. This institutional design is guided by the policy goal of providing clear legal certainty for Japan's AI industry.

3.1.2. Strict model: China's pre-examination of legitimate sources

In contrast to Japan's lenient stance, China's institutional response centers on prior compliance review. Article 7 of the *Interim Measures for the Administration of Generative Artificial Intelligence Services* requires that the sources of training data be legitimate, establishing "legitimate source" as a prerequisite for data training. In the regulatory approach to preventing infringement risks of training data, the primary link is to ensure the legitimate sources of training data at the input stage. The ethical core of this model reflects a strong protection orientation for copyright holders' rights.

3.1.3. Balanced model: transformative use and opt-out mechanisms in the United States and the United Kingdom

The institutional responses of the United States and the United Kingdom represent efforts to strike a middle ground between industrial promotion and rights protection, but their balancing mechanisms differ significantly.

U.S. regulation of AI training data use relies on the four-factor framework for fair use, with the concept of transformative use at its core. Judging from the 2025 cases of Bartz v. Anthropic and Kadrey v. Meta, U.S. judicial practice recognizes the reasonableness of AI training data use in individual cases, but rulings are highly dependent on specific circumstances, with obvious case-by-case adjudication characteristics. Notably, judges' rulings generally focus on potential market impacts, leaving room for adjustment in future rulings at the legal level.

A key theoretical detail to clarify here is the negating effect of competitive substitutability on fair use. If the output of an AI model constitutes a substantial substitute for the market of the original work, even if the training has transformative characteristics, the fourth factor of fair use (impact on the potential market of the work) may still negate the overall finding. There are two core legal challenges in the use of AI training data: the legality of using copyrighted works as training data, and the copyright of content derived from copyrighted training data. This judicial model of "coexisting limited exemptions and systematic uncertainty" reflects the U.S. system's preference for ex post case-by-case weighing at the ethical level. Its advantage is the flexibility to adapt to technological development, but the cost is severe uncertainty in behavioral expectations for both creators and developers.

The European Union's institutional framework centers on the Text and Data Mining (TDM) exception rules in the *Copyright Directive on the Digital Single Market* (CDSM Directive). The EU

DOI: 10.54254/2753-7048/2026.LD33864

adopts a dual-track system for TDM exceptions, distinguishing between scientific research purposes and general purposes. It stipulates that research organizations and cultural heritage institutions may conduct TDM for scientific research purposes, while TDM for general purposes is subject to reservation by right holders [10]. In its institutional evolution after Brexit, the United Kingdom has continued this opt-out logic.

3.2. Analysis of ethical deficiencies in regulation across jurisdictions

There are significant differences in the requirements for the legitimacy of training data sources across jurisdictions, which reflect varying degrees of moral risks. Japan's "non-appreciation purpose" exception theoretically imposes no independent legitimacy constraint on data sources. When reproduction in machine learning is classified as non-appreciative use, the legitimacy of data sources may not constitute an independent infringement element under the law, raising ethical concerns about the inclusion of data from illegal websites in training sets. In contrast, the EU's TDM exception is premised on "lawful access", while China explicitly regards legitimate sources as a prerequisite for prior review.

The moral risk regarding the legitimacy of sources stems from inconsistent criteria for defining "legitimate" across jurisdictions. Under Japanese law, "legitimate" refers only to the legality of the act of non-appreciative use itself, not the legality of the data acquisition channel; under Chinese law, "legitimate sources" cover both legal acquisition channels and legal use acts. This difference creates regulatory arbitrage space in the globalized supply chain of AI training data, where cross-border data flows are almost unrestricted. Another common ethical deficiency in the regulation of training data use across jurisdictions is the lack of an institutional design to systematically feed back the benefits of data use to creators. Japan's non-appreciative use exception, the U.S. transformative use defense, and the EU's TDM exception are all centered on exemption or restriction. They either exempt users from payment obligations or grant right holders an opt-out right without guaranteeing their right to remuneration.

4. Copyrightability of generated content and ethics of authorship

As generative AI becomes deeply involved in the creative process, the traditional human-author-centered legitimacy basis of copyright is called into question. None of the three classic theories—the utilitarian incentive theory, the labor desert theory, and the personality theory—can provide a coherent justification for AI-generated content [11]. Around this core theoretical dilemma, countries around the world have developed multiple response paths. The evolution of paradigms from the Absolute Control Doctrine to the Substantial Contribution Doctrine, and the ethical expansion from the protection of signature rights to the risk of identity forgery, constitute an ongoing and unresolved institutional dialogue.

4.1. Absolute control doctrine: a conservative stance on the dignity of human creation

The core of the Absolute Control Doctrine is that a necessary condition for copyright protection is human absolute control over the form of expression of a work, with AI acting merely as a tool. If AI exercises autonomy in generating the form of expression, copyrightability is excluded. The stance of the U.S. Copyright Office (USCO) in the Thaler case and related rulings is a typical embodiment of this paradigm.

DOI: 10.54254/2753-7048/2026.LD33864

4.2. Linear modification doctrine: a gradual shift in creative intent

The "linear modification" proposal is an important adjustment within the Absolute Control paradigm. It argues that if human modification of the initial AI output reaches a sufficient level of creativity, it may constitute copyrightable expression. From an ethical perspective, it attempts to strike a balance between two opposing values: maintaining the human-centric nature of copyright and acknowledging the reality of technology-assisted creation. However, its theoretical vulnerability lies in its underlying assumption. In actual workflows, human intervention and AI generation are highly iterative and non-linear, and simplifying modification to linear contribution fails to capture the true complexity of the creative process.

4.3. Substantial contribution doctrine: an inclusive paradigm for the creative process

In contrast to the Absolute Control Doctrine, the Substantial Contribution Doctrine holds that the criterion for copyrightability should shift from "whether humans exercise absolute control over AI" to "whether humans make substantial original intellectual contributions in creation". It does not presuppose a specific human-machine interaction model, but uses the density of creative contribution as the judgment criterion. Based on the level of human control and contribution in the AI generation process, it introduces dimensions such as technical controllability and the density of human intent, classifying generative AI into three tiers: strong protection, weak protection and no protection. The core divergence among the three paradigms lies not in whether AI can be an author, but in how to set the threshold for sufficient human participation. The Absolute Control Doctrine has the highest threshold, the Substantial Contribution Doctrine the lowest, and the Linear Modification Doctrine in between. This difference in thresholds stems from different understandings of the purpose of the copyright system: the former emphasizes the moral rights of copyright, while the latter focuses on economic rights. From a risk perspective, copyright risks run through the entire life cycle of AI, affecting multiple stakeholders including copyright owners of training data, AI users, and developers and operators of training models. This indicates that a single paradigm cannot cover all interest dimensions, and paradigm selection involves value trade-offs.

5. Ethical reconstruction of liability allocation

5.1. Core liability of developers

In AIGC copyright infringement scenarios, developers are at the core of the liability chain. When developers fail to take reasonable measures despite knowing that training data contains copyrighted content, they should bear a higher degree of legal liability. Copyright infringement patterns of generative AI are more complex than traditional online infringement, and infringing subjects cannot be simply defined as direct or indirect infringers [12].

From a comparative law perspective, drawing on the strict liability system in the nuclear energy field, some scholars advocate imposing limited, strict and exclusive third-party liability on developers of cutting-edge AI models, especially when their systems cause substantial harm [13]. In high-risk fields such as medical AI, developers may be held strictly liable for design defects or failure to warn of specific limitations of the system. Japan's "presumption of dependency" rule also reflects this logic: when developers know or should know about the infringement risks of training data, they are presumed to have a dependent relationship with the infringement consequences of the generated content, thus bearing joint and several liability.

DOI: 10.54254/2753-7048/2026.LD33864

5.2. Duty of care of platforms

In AIGC scenarios, as providers of generative AI services and distribution channels, platforms constitute contributory infringement and bear joint and several liability if they know or should know that users use platform tools to generate infringing content but fail to take reasonable measures. The determination of copyright infringement liability for AI service providers should not simply adopt the traditional "notice-and-takedown" rule for online service providers, but establish a more active duty of care based on their degree of control over the generation process and technical capabilities. In addition, this paper proposes using technical means such as blockchain to manage the full life cycle of AIGC products, enhancing the traceability of copyright protection [14].

5.3. Ethical insurance and risk-sharing mechanisms

A compulsory insurance system is widely discussed in AI risk governance. Drawing on precedents in the nuclear energy industry, it is recommended to implement compulsory liability insurance for developers of cutting-edge AI. Through the quasi-regulatory function of insurance companies, ex ante risk control can be realized. Insurers will be involved in causal risk modeling, continuous monitoring, promoting industry norms and providing loss prevention guidance. In practice, a three-tier insurance structure can be established: (1) developers purchase private liability insurance compulsorily; (2) an industry-managed risk pool absorbs routine non-catastrophic losses; (3) a government-backed reinsurance mechanism addresses tail risks. For uninsurable risks beyond the capacity of commercial insurance, the state can act as the insurer of last resort.

6. Toward collaborative governance: an ethics-embedded intellectual property framework

For the use of massive data in base model training, it is recommended to introduce a system combining statutory licenses and collective management. AI developers can obtain the right to use specific categories of works for base model training by paying statutory royalty rates to a designated collective management organization; right holders may opt out of collective management, but will bear the consequence of increased transaction costs. This scheme has obvious advantages: it significantly reduces transaction costs and solves the problem of mass authorization; it ensures fair compensation for creators, correcting the ethical unfairness of Japan's "free use" model; it preserves the right holders' freedom of choice and respects their right to self-determination.

For data use in the fine-tuning stage, separate authorization is required. Fine-tuning optimizes models for specific fields or styles, and the selection of training data determines the characteristics and quality of output content. Separate authorization can prevent disputes over style imitation. If AI developers intend to make models imitate the style of a specific artist, they need to obtain authorization from that artist.

A mandatory disclosure obligation for AI use shall be introduced into the copyright registration system. When submitting works for registration, applicants must declare whether AI tools were used to assist creation, the specific tools and their functions, and the specific content and scope of human contributions. The 2025 report of the U.S. Copyright Office provides an actionable reference template for this practice.

Given the institutional divergences on AI intellectual property issues across countries, it is difficult to reach a binding international treaty in the short term. A more realistic path is to explore the possibility of soft law guidelines on "AI Ethics and IP" within the framework of WIPO, providing a reference framework for national legislation and judicial practice. The core content of

DOI: 10.54254/2753-7048/2026.LD33864

the guidelines should include: minimum standards for the transparency principle, diversified mechanisms for fair compensation, common benchmarks for copyrightability, and the proportionality principle for liability allocation. The soft law guidelines shall establish a regular evaluation mechanism to periodically assess the implementation effects of the guidelines and technological developments, identify emerging ethical and legal challenges, and propose revision suggestions. At the same time, a technical response mechanism shall be established to initiate research and dialogue in a timely manner in response to major technological changes.

7. Conclusion

As artificial intelligence evolves from an auxiliary tool to an autonomous decision-making system, the global legal system is shifting from regarding it as an exception of neutral technology to systematically integrating ethical dimensions. In the current intellectual property field, disputes over the copyright of training data and the determination of authorship of AI-generated content have fallen into a dilemma due to the human-centrism of traditional copyright law, requiring refined ethical and legal analytical tools to balance innovation and rights. The rise of autonomous agents poses even deeper challenges. The opacity and unpredictability of such systems, which can perceive, learn and make decisions independently, render traditional liability determination mechanisms based on human fault ineffective, creating fundamental difficulties in liability attribution and rights allocation.

"Justice is the first virtue of social institutions, as truth is of systems of thought." [15] In the face of these challenges, research on intellectual property law and ethics must break through the limitations of a single discipline and move toward an interdisciplinary collaborative research paradigm. Legal studies must integrate computer science to understand technical boundaries, combine ethics to respond to algorithmic justice, and collaborate with social sciences to assess institutional impacts. Only in this way can we provide theoretical reserves for legal construction in the era of autonomous agents, both adapting to technological changes and adhering to the value orientation of safeguarding human well-being.

References

[1] John Locke, Two Treatises of Government, (Peter Laslett ed.) Cambridge University Press, 1988, § 27 (1689).
[2] Immanuel Kant, Groundwork of the Metaphysics of Morals 38 (Mary Gregor trans., Cambridge University Press 1998) (1785).
[3] Jyh-An Lee, Reto Hilty & Kung-Chung Liu, Roadmap to Artificial Intelligence and Intellectual Property: An Introduction, in Artificial Intelligence and Intellectual Property 1 (Jyh-An Lee et al. eds., 2021).
[4] I Kadek Widi Pranajaya & I Gede Agus Kurniawan, Intellectual Property Rights in the Age of AI and Big Data: A Reinterpretation of the Concept of Creativity in Legal Philosophy, 24 Pena Justisia 7042 (2025).
[5] William M. Landes & Richard A. Posner, An Economic Analysis of Copyright Law, The Journal of Legal Studies, vol. 18, no. 2, pp. 326–333 (1989).
[6] Japanese Copyright Act (Act No. 48 of 1975), Article 2, Paragraph 1, Item 1.
[7] Liu Zubing, Legal Paths, Obstructions and Institutional Improvement for Generative Artificial Intelligence's Use of Pre-Existing Works as Training Data, Journal of Xihua University (Philosophy and Social Sciences), Vol. 44, No. 2, pp. 19–20 (2025).
[8] Japanese Copyright Act (Revised 2018), Article 47-5.
[9] Li Kexin & Xiao Dongmei, Japan's Fair Use Rules for Generative Artificial Intelligence Training Data and Their Enlightenment, Library Forum, Vol. 45, No. 9, pp. 93–101 (2025).
[10] Jiao Heping, Copyright Risks and Resolution Paths for Data Acquisition and Utilization in AI Creation, Contemporary Law Review, No. 4, p. 135 (2022).

DOI: 10.54254/2753-7048/2026.LD33864

[11] Ezieddin Elmahjub, The Algorithmic Muse and the Public Domain: Why Copyright's Legal Philosophy Precludes Protection for Generative AI Outputs, 58 Computer Law & Security Review 106170, 106170 (2025).
[12] Cristian Trout, Liability and Insurance for Catastrophic Losses: The Nuclear Power Precedent and Lessons for AI, arXiv: 2409.06673 (2024).
[13] Yinqiu Liu et al., Blockchain-Empowered Lifecycle Management for AI-Generated Content Products in Edge Networks, 31 IEEE Wireless Communications 286, 286 (2024).
[14] Nicholas Stetler, Reinsuring AI: Energy, Agriculture, Finance & Medicine as Precedents for Scalable Governance of Frontier Artificial Intelligence, 50 Southern Illinois University Law Journal 189, 32 (2025).
[15] John Rawls, A Theory of Justice 3 (rev. ed. 1999).