

The Impact of Generative Artificial Intelligence on Copyright Law and Intellectual Property Protection: A Comparative Doctrinal Analysis

Rafaqat Ali*

School of Law, Northwest University, Xi'an, China

Email: rafaqatalishah.112@gmail.com

Faraz Jan

Department of Law, University of Lincoln, Lincoln, United Kingdom

Email: Farazjan22033@gmail.com

Kanwal Zahara

Department of Law, Fatima Jinnah Women University, Rawalpindi, Pakistan

Email: kanwallb786@gmail.com

Hasnain Haider

Department of Computer Science, University of Baltistan, Skardu, Pakistan

Email: hassubalti7@gmail.com

ABSTRACT

Generative artificial intelligence has placed copyright law under a kind of pressure it was not designed to absorb. Two questions sit at the centre of the problem. The first concerns inputs: whether copying millions of protected works to train a model is an infringing act or a lawful, even socially valuable, form of analysis. The second concerns outputs: whether text, images or music produced by a model can be owned at all, and by whom. This article examines both questions through a doctrinal and comparative study of the United States, the European Union, the United Kingdom and China, with briefer reference to Japan. It draws on statutes, regulatory guidance and the first wave of judicial decisions, including *Thaler v. Perlmutter*, *Bartz v. Anthropic*, *Kadrey v. Meta*, *Getty Images v. Stability AI* in the English High Court, *GEMA v. OpenAI* in Munich and *Li v. Liu* in the Beijing Internet Court. The analysis finds that courts now broadly agree that works lacking human creative control are unprotected in the United States, while the United Kingdom and China reach more permissive results by different routes. On training, the picture is less settled. American courts have so far treated training on lawfully acquired works as transformative, yet the Munich court held that a model which memorises lyrics itself contains a reproduction. Patent law is more uniform: every major office has refused to name an AI system as an inventor. The article argues that litigation alone cannot resolve these tensions and proposes a three-part reform built on training-data transparency, statutory remuneration for rights holders and a short neighbouring right for AI-generated output.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Keywords: Generative Artificial Intelligence; Copyright; Authorship; Fair Use; Text And Data Mining; AI Inventorship; EU AI Act; Comparative Law

Introduction

When OpenAI opened ChatGPT to the public in late November 2022, few lawyers expected the product to become a test case for the foundations of intellectual property. Within about two months the service was reported to have reached roughly one hundred million monthly users (Hu, 2023). Image generators such as Midjourney and Stable Diffusion had arrived a few months earlier, and music and video systems followed. What these tools share is a method. They are trained on very large collections of human-made material, most of it protected by copyright, and they then produce new material that is statistically consistent with what they have seen. That method touches copyright twice: once when the training data is copied, and again when the output is released into the world.

The scale of the copying is worth stating plainly, because it shapes the legal analysis. GPT-3 was trained on a corpus built from about 45 terabytes of compressed Common Crawl text, filtered down to around 570 gigabytes, together with book collections and Wikipedia (Brown et al., 2020). LAION-5B, the open dataset used for several image models, lists some 5.85 billion image and text pairs (Schuhmann et al., 2022). No licence was sought for most of this material. Developers say this is analysis, not exploitation, and compare it to a person reading widely. Authors, photographers, newspapers and music publishers reply that their work has been taken to build products that now compete with them.

The output side raises a different difficulty. Copyright has long assumed a human author. The United States Copyright Office has refused to register works produced without meaningful human creative control, and the federal courts have agreed (*Thaler v. Perlmutter*, 2023, 2025). The Supreme Court declined to review that position in March 2026. The consequence is that a company may spend heavily to generate images or text that anyone can copy. The United Kingdom, by contrast, has since 1988 recognised a category of computer-generated works owned by the person who makes the necessary arrangements, and a Chinese court has granted protection to an AI-assisted image on the basis of the user's creative choices in prompting.

Patent law adds a third strand. Dr Stephen Thaler's applications naming the DABUS system as inventor were rejected in the United States, the United Kingdom, before the European Patent Office and, after an initial success, in Australia. These decisions rested on statutory wording rather than on any considered view of whether AI-generated inventions ought to be patentable, which leaves the policy question open.

This article asks how far existing intellectual property doctrine can accommodate generative AI, where it fails, and what reform would best balance the interests of creators, developers and the public. Section 2 reviews the literature. Section 3 explains the method. Sections 4 to 6 analyse authorship, training data and inventorship. Section 7 compares the main jurisdictions. Section 8 sets out reform proposals, and Section 9 concludes.

Literature Review

Early work on computer-generated works

Scholarly concern with machine creativity predates today's models by some years. Bridy (2012) was among the first to argue that algorithmic authorship would eventually force copyright law either to stretch the concept of the author or to accept that some commercially valuable works would fall into the public domain. She suggested an analogy with works made for hire, which would place ownership with the party that commissions or controls the machine. Yanisky-Ravid (2017) developed a related argument around accountability, proposing that whoever deploys an AI system should bear both the benefits and the liabilities of its output. Hristov (2017) reached a similar conclusion, favouring ownership by the programmer or user over any recognition of the machine itself.

Comparative work soon followed. Guadamuz (2017) examined how originality was defined in several legal systems and concluded that the United Kingdom's computer-generated works provision was the only developed legislative answer, though one resting on an originality standard that sat uneasily with European case law. Gervais (2020) pressed the opposite view, arguing that copyright should protect only works resulting from human creative choices and that extending it to machine output would weaken the justification for the whole system.

Authorship and originality

The most detailed doctrinal treatment of authorship is that of Ginsburg and Budiardjo (2019). They traced the human authorship requirement through American case law and concluded that it could not easily be abandoned without legislation. They also proposed a useful distinction between works in which the machine is a tool of a human author and works in which the human's contribution is too remote to count as authorship. That distinction now underpins the Copyright Office's registration guidance (U.S. Copyright Office, 2023, 2025a). The monkey selfie litigation, *Naruto v. Slater* (2018), is often cited in this literature, although it was decided on standing and not on the merits of non-human authorship.

European scholarship frames the same issue through the Court of Justice's originality test. A work is protected where it is the author's own intellectual creation (*Infopaq*, 2009), reflecting free and creative choices that stamp the work with the author's personal touch (*Painer*, 2011). Commentators generally accept that a purely machine-produced work cannot satisfy that test, while disagreeing about how much prompting or editing is enough to bring an output back within it (Quintais, 2025).

Training data and fair use

The literature on training data grew quickly after 2020. Lemley and Casey (2021) argued that machine learning should generally qualify as fair use because models learn unprotected facts and patterns rather than expression, an argument they called fair learning. Sag (2023) broadly agreed but identified memorisation as a real risk, noting that a model able to reproduce its training data verbatim is harder to defend. Henderson et al. (2023), writing from a combined legal and technical perspective, showed that foundation models can regurgitate protected text and argued that developers should build technical mitigations if they wish to rely on fair use.

Samuelson (2023) took a more cautious position, suggesting that outcomes would depend heavily on the fourth fair use factor and on whether courts accept that a licensing market for training data exists.

Empirical computer science has given this debate a factual base. Carlini et al. (2021) extracted verbatim training sequences from GPT-2, and Carlini et al. (2023) recovered near-copies of training images from diffusion models. Lee, Cooper and Grimmelmann (2023) argued that copyright analysis must follow the whole generative AI supply chain, from dataset creation to deployment, because liability may attach at several points and to several actors. In Europe, Senftleben (2023) proposed that authors be compensated through a levy on AI output rather than through a consent-based system, on the ground that individual licensing at this scale is unworkable.

Patent law and the research gap

Abbott (2016) anticipated the inventorship problem before DABUS was filed, arguing that denying patents for machine-generated inventions would reduce incentives to develop inventive AI. His later book develops a principle of AI legal neutrality, under which the law should not treat an act differently merely because a machine performed it (Abbott, 2020). WIPO (2020) summarised the international debate and identified inventorship, ownership and training-data use as open questions for member states.

What remains missing is work that places the copyright input problem, the copyright output problem and the patent inventorship problem side by side, and that takes account of the judgments delivered in 2025 and early 2026. Most existing studies were written before any court had decided a generative AI training dispute on the merits. This article addresses that gap.

Research Methodology

The study uses doctrinal legal research combined with functional comparative analysis. The doctrinal component identifies and interprets the primary sources: statutes, regulations, official guidance and judicial decisions (Hutchinson & Duncan, 2012). The comparative component follows the functional method, which asks how different legal systems respond to the same practical problem rather than comparing rules that merely share a label (Zweigert & Kötz, 1998).

Four jurisdictions were selected. The United States hosts most leading AI developers and the largest body of litigation. The European Union has adopted the most comprehensive regulatory instrument, the AI Act. The United Kingdom offers a distinctive statutory model for computer-generated works. China combines rapid judicial development with explicit administrative rules on generative AI. Japan is referred to where its broad training exception offers a useful contrast. Sources were collected up to June 2026. Where cases are still pending, the article describes the issues raised rather than predicting outcomes. Because the study is doctrinal, it does not rely on quantitative data; the figures and tables that follow summarise legal positions and are not statistical results.

Authorship and the Protection of AI Output

The human authorship requirement in the United States

American copyright law has tied protection to human creativity since at least *Burrow-*

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Giles Lithographic Co. v. Sarony (1884), where the Supreme Court upheld copyright in a photograph because the photographer had posed the subject, arranged the lighting and chosen the costume. The camera was treated as an instrument of the photographer's mind. *Feist Publications v. Rural Telephone Service* (1991) later confirmed that originality requires independent creation and a modicum of creativity. Neither case contemplated a tool that makes the expressive choices itself.

That question reached the courts in *Thaler v. Perlmutter*. Thaler sought to register an image titled *A Recent Entrance to Paradise*, which he said had been created autonomously by his Creativity Machine, and he disclaimed any human contribution. The district court upheld the refusal, holding that human authorship is a bedrock requirement of copyright (*Thaler v. Perlmutter*, 2023). The D.C. Circuit affirmed in March 2025, reading the Copyright Act as a whole, including its provisions on duration, inheritance and signatures, as presupposing a human author (*Thaler v. Perlmutter*, 2025). The Supreme Court denied certiorari on 2 March 2026. The fully autonomous case is therefore settled in the United States, at least until Congress acts.

The harder cases involve mixed authorship. In the *Zarya of the Dawn* decision, the Copyright Office accepted registration of a comic book's text and the selection and arrangement of its images, but refused protection for the individual images produced by Midjourney, reasoning that the author's prompts did not control the specific visual result (U.S. Copyright Office, 2023a). The Office's formal guidance, issued the following month, requires applicants to disclose AI-generated material and to claim only their own contribution (U.S. Copyright Office, 2023b). Its 2025 report on copyrightability went further, concluding that prompts alone generally do not give a user enough control over the output to count as authorship, though expressive inputs, such as an uploaded drawing that the output preserves, and later creative editing may be protected (U.S. Copyright Office, 2025a). The practical line, then, is drawn by control rather than effort. A user who spends hours refining prompts may still own nothing if the system decides how the result looks.



Figure 1. The spectrum of human contribution and the treatment of output in three jurisdictions. Protection is clearest at the left of the spectrum; the jurisdictions diverge most sharply at the right. Prepared by the authors from the sources discussed in Sections 4 and 7.

The United Kingdom's computer-generated works provision

The United Kingdom chose a different route in the Copyright, Designs and Patents Act 1988. Section 9(3) provides that where a literary, dramatic, musical or artistic work is computer-generated, the author is taken to be the person by whom the arrangements necessary for its creation are undertaken. Section 178 defines a computer-generated work as one generated in circumstances where there is no human author, and section 12(7) limits protection to fifty years from the year of creation. The provision has been applied rarely. In *Nova Productions v. Mazooma Games* (2007), the Court of Appeal held that the frames displayed by a video game were computer-generated and that the programmer, not the player, had made the necessary arrangements.

The provision has two weaknesses. It assumes originality in works that by definition lack a human author, and it is hard to reconcile with the intellectual creation standard that shaped United Kingdom law while it was bound by European case law (Guadamuz, 2017). It is also unclear who makes the arrangements when a user prompts a model built by someone else. The developer supplies the system; the user supplies the prompt. The government's December 2024 consultation asked whether section 9(3) should be kept, reformed or repealed, which suggests that even its supporters doubt whether a 1988 rule fits present technology (UK Intellectual Property Office, 2024).

China's creative-choice approach

Chinese courts have been more willing to protect AI-involved works. In 2019 the Shenzhen Nanshan District People's Court held that a financial report produced with Tencent's Dreamwriter software was protected, because the selection and arrangement of data, templates and parameters reflected the choices of Tencent's staff (*Shenzhen Tencent v. Shanghai Yingxun*, 2019). In *Li v. Liu* (2023), the Beijing Internet Court went further, recognising copyright in an image generated with Stable Diffusion. The court emphasised that the plaintiff had chosen the model, written and revised detailed prompts, adjusted parameters and selected among outputs, and it treated these steps as reflecting his aesthetic judgement. In substance this is the same inquiry the Copyright Office conducts, but it reaches a more generous answer on similar facts. That difference is less about doctrine than about how much control a court is willing to find in the act of prompting.

Training Data, Memorisation and the Limits of Existing Exceptions

Where copying occurs

Copying can occur at several stages of the generative AI pipeline, and each stage raises its own legal question (Lee et al., 2023). Figure 2 maps these stages. Works are copied when datasets are assembled, again when they are loaded for training, and arguably a third time if the trained model retains them in a form that can be recovered. Outputs may then reproduce or closely resemble particular works. Keeping these stages separate matters, because an exception may cover one stage and not another, and because different parties may be responsible for different stages.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Fair use in the United States

Section 107 of the Copyright Act asks courts to weigh the purpose and character of the use, the nature of the work, the amount taken and the effect on the market. After *Campbell v. Acuff-Rose Music* (1994), the first factor turns largely on whether the use is transformative. *Authors Guild v. Google* (2015) held that scanning millions of books to create a searchable index was transformative because it served a different function and did not supply substitutes for the books. Developers rely heavily on that decision. *Andy Warhol Foundation v. Goldsmith* (2023), however, narrowed the transformative use inquiry by focusing on whether the specific use shares the purpose of the original, which may matter where AI output competes with the works used to train it.

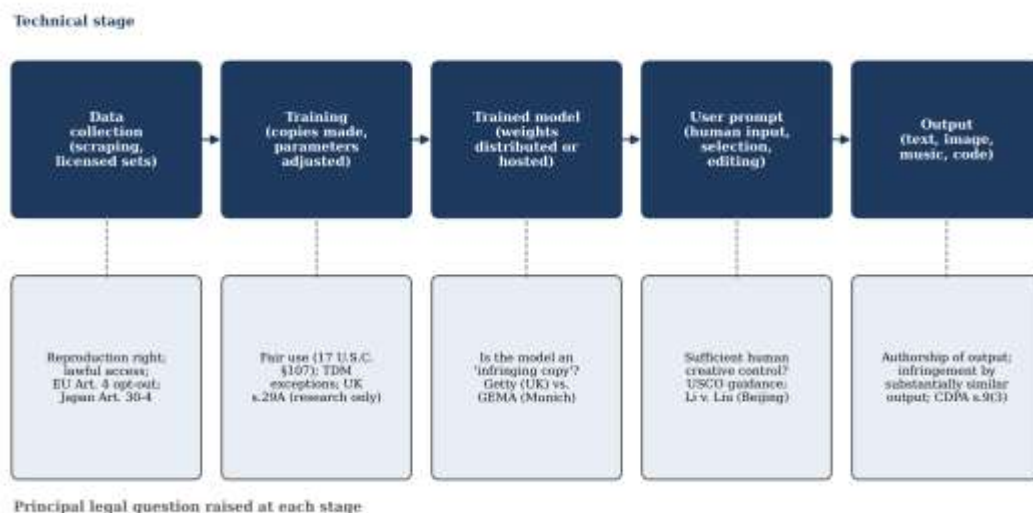


Figure 2. Stages of the generative AI lifecycle and the principal legal question raised at each. Prepared by the authors.

The first merits decisions arrived in 2025 and did not all point the same way. In *Thomson Reuters v. Ross Intelligence* (2025), a case about a non-generative legal research tool, the court rejected fair use because Ross had copied Westlaw headnotes to build a competing product. In *Bartz v. Anthropic* (2025), Judge Alsup held that training a language model on lawfully purchased books was highly transformative and fair, but that downloading and keeping millions of pirated books to build a central library was not. The case later settled for a reported 1.5 billion US dollars. Two days after *Bartz*, *Kadrey v. Meta Platforms* (2025) granted summary judgment to Meta, but Judge Chhabria stressed that the authors had lost because they had not proved market harm, and he suggested that evidence of market dilution by AI-generated substitutes could succeed in a future case. The Copyright Office's pre-publication report on training reached a similar middle position: some training will be fair use, but uses that produce expressive substitutes for the training works in their own markets are unlikely to be (U.S. Copyright Office, 2025b).

Two lessons follow. First, how the data was obtained matters as much as what the model does with it. Second, the fourth factor, market effect, is where the most important fights will take place. The emerging licensing market, including agreements between OpenAI and news publishers such as the Associated Press, gives rights

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

holders a stronger argument that unlicensed training causes cognisable harm (Samuelson, 2023).

Europe: text and data mining and the memorisation problem

The European Union answered the training question in advance, if incompletely. Articles 3 and 4 of the Digital Single Market Directive (Directive (EU) 2019/790) create two exceptions for text and data mining. Article 3 covers research organisations and cultural heritage institutions. Article 4 covers everyone else but allows rights holders to reserve their rights, for online content by machine-readable means. The AI Act (Regulation (EU) 2024/1689) links to this framework. Article 53(1)(c) requires providers of general-purpose AI models to adopt a policy to comply with Union copyright law, including reservations under Article 4, and Article 53(1)(d) requires a publicly available summary of training content. These obligations have applied since 2 August 2025.

The Munich Regional Court's decision in *GEMA v. OpenAI* (2025) shows where this framework runs out. GEMA demonstrated that GPT-4 and GPT-4o could reproduce the lyrics of nine German songs in response to simple prompts. The court accepted that copying to create a training corpus could fall within the text and data mining exception, but held that the memorisation of the lyrics within the model's parameters was itself a reproduction not covered by the exception, and that OpenAI, not its users, was responsible for infringing outputs. OpenAI has appealed. The ruling matters because it treats the model as capable of containing the works on which it was trained, a proposition that developers have long denied.

The English High Court took a narrower view on different evidence. In *Getty Images v. Stability AI* (2025), Getty abandoned its training claim because the training had taken place outside the United Kingdom. On the remaining secondary infringement claim, Joanna Smith J accepted that a model can be an article for the purposes of the Act, but held that Stable Diffusion's weights were not an infringing copy because they did not store Getty's images. Getty succeeded only on limited trade mark claims concerning watermark-like features in early outputs. Read together, *Getty* and *GEMA* do not truly conflict. Each turns on what the evidence showed about memorisation, and each suggests that the capacity of a particular model to regenerate particular works will decide liability.

Table 1. Selected generative AI copyright disputes and their status (as at June 2026)

Case	Forum	Core issue	Outcome or status
<i>Thomson Reuters v. Ross Intelligence</i>	D. Del. (US)	Copying headnotes to train a legal AI tool	Fair use rejected (Feb 2025); interlocutory appeal
<i>Bartz v. Anthropic</i>	N.D. Cal. (US)	Training on purchased and pirated books	Training fair; pirated library not (Jun 2025); settled
<i>Kadrey v. Meta Platforms</i>	N.D. Cal. (US)	Training Llama on books	Summary judgment for Meta for want of market-harm evidence (Jun 2025)
<i>The New York Times v. Microsoft &</i>	S.D.N.Y. (US)	Training and verbatim output of news articles	Pending

Case	Forum	Core issue	Outcome or status
<i>OpenAI</i>			
<i>Andersen v. Stability AI</i>	N.D. Cal. (US)	Image-model training on artists' works	Pending
<i>Getty Images v. Stability AI</i>	High Court (UK)	Secondary infringement; trade marks	Model weights not an infringing copy; limited trade mark infringement (Nov 2025)
<i>GEMA v. OpenAI</i>	LG München I (DE)	Memorisation of song lyrics	Infringement found; TDM exception inapplicable to memorisation (Nov 2025); on appeal

Other approaches

Japan offers the most permissive model. Article 30-4 of its Copyright Act, introduced in 2018, permits the use of works for data analysis and other purposes that do not involve enjoying the expression, subject to a proviso against unreasonably prejudicing the rights holder. Official guidance issued in 2024 clarified that the exception does not cover training intended to reproduce the expression of particular works (Agency for Cultural Affairs, 2024). The United Kingdom's own exception, section 29A of the 1988 Act, is limited to non-commercial research, and the government's 2024 consultation proposed a broader exception with a rights-reservation mechanism modelled on the European approach (UK Intellectual Property Office, 2024). China's Interim Measures for generative AI services require providers to use training data from lawful sources and not to infringe intellectual property rights, though they say little about how that duty is to be enforced.

Patent Law and AI Inventorship

Patent statutes are more explicit than copyright statutes about who may create. The United States Patent Act defines the inventor as the individual or individuals who invented the subject matter (35 U.S.C. § 100(f)), and its oath requirements refer to the inventor's own belief. The European Patent Convention requires the designation of the inventor under Article 81. The DABUS applications tested these rules in parallel, and Table 2 summarises the results.

Table 2. Outcomes of the DABUS applications naming an AI system as inventor.

Jurisdiction	Decision	Result and reasoning
United States	<i>Thaler v. Vidal</i> (Fed. Cir. 2022); cert. denied 2023	Refused: 'individual' in the Patent Act means a natural person
European Patent Office	J 8/20 and J 9/20 (Legal Board of Appeal, 2021)	Refused: the designated inventor must be a person with legal capacity
United Kingdom	<i>Thaler v. Comptroller-General</i> [2023] UKSC 49	Refused: inventor must be a natural person; ownership of the machine gives no title
Australia	<i>Commissioner of Patents v.</i>	First-instance acceptance reversed; special

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Jurisdiction	Decision	Result and reasoning
	<i>Thaler</i> [2022] FCAFC 62	leave to the High Court refused
Germany	BGH, X ZB 5/22 (2024)	A natural person must be named, but a statement that the invention was produced by AI may be added
South Africa	Patent ZA 2021/03242 (2021)	Granted without substantive examination; limited precedential value

The consensus is clear but narrow. These rulings decide who may be named on an application, not whether an invention developed with AI can be patented. The practical question is how much a human must contribute. In February 2024 the USPTO issued guidance applying the joint inventorship factors from *Pannu v. Iolab* to human contributions made with AI assistance. In November 2025 the Office rescinded that guidance and replaced it with a simpler position: AI systems are tools, there is no separate standard for AI-assisted inventions, and the ordinary conception test applies to the human inventor (U.S. Patent and Trademark Office, 2025). The change makes it easier for applicants to name a human who directed an AI tool, but it leaves open how conception should be assessed when a system proposes thousands of candidate molecules and a scientist selects one. The German Federal Court of Justice's approach, which allows the application to state that AI produced the invention while still naming a human, is a sensible compromise that preserves transparency without inventing a new legal person.

Comparative Analysis

The four jurisdictions reach different answers because they start from different institutional habits. The United States develops the law mainly through litigation and agency guidance. Congress has considered several bills, including proposals on training-data disclosure and on digital replicas of voice and likeness, but none has been enacted. The advantage is careful, fact-based reasoning. The disadvantage is time, since large copyright cases take years while model generations turn over in months.

The European Union has preferred ex ante regulation. The AI Act does not change substantive copyright law, but by requiring a copyright compliance policy and a public summary of training data, it gives rights holders the information they need to enforce their reservations. Its weakness is that the underlying questions of authorship and memorisation are left to national courts and, eventually, to the Court of Justice.

The United Kingdom has the only developed statutory rule for machine output but has hesitated on training. A proposal to extend the text and data mining exception to commercial use was dropped in 2023 after objections from the creative industries, and the 2024 consultation reopened the question. China has combined permissive courts with prescriptive administrative rules, recognising AI-assisted works while imposing lawful-source and labelling duties on providers. Table 1 and Table 3 set out the main differences, and Figure 3 shows how quickly the key decisions have accumulated.

Online ISSN

3007-3197

Print ISSN

3007-3189


<http://amresearchreview.com/index.php/Journal/about>

Figure 3. Timeline of selected judicial and regulatory milestones on generative AI and intellectual property, 2019 to 2026. Prepared by the authors.

Table 3. Comparative summary of legal positions on generative AI and intellectual property.

Issue	United States	European Union	United Kingdom	China
Protection of purely AI-generated output	No (<i>Thaler</i>)	No (own intellectual creation required)	Yes, 50 years (CDPA s. 9(3))	Unsettled
AI-assisted output	Human elements only	Where creative choices are shown	Yes	Yes where prompting reflects aesthetic choice (<i>Li v. Liu</i>)
Legal basis for training	Fair use, case by case	TDM exceptions (DSM Arts. 3–4) with opt-out	Non-commercial research only (s. 29A)	Lawful-source duty (Interim Measures)
Training transparency duty	None enacted	Yes (AI Act Art. 53)	Under consultation	Partial (filing and labelling rules)
AI as patent inventor	No	No (EPO)	No	No

Two broader points emerge. The first is convergence on outcomes that matter commercially. Every jurisdiction protects works in which a human makes substantial creative choices, and every patent office refuses to treat a machine as an inventor. The second is divergence on the training question, which is the one with the greatest economic weight. A developer may train lawfully in Japan, be sued for the same model in Germany and face a fair use trial in California. Divergence of this kind invites regulatory arbitrage, since training can be moved to the most permissive jurisdiction, as the jurisdictional collapse of Getty's training claim in London illustrates.

Proposals for Reform

Three reforms would address the gaps identified above without halting the development of a useful technology. They are summarised in Figure 4.




Figure 4. A three-pillar framework for reform, supported by coordinated minimum standards. Prepared by the authors.

The first is training-data transparency. Rights holders cannot enforce reservations or claim payment if they do not know whether their work was used. The AI Act's public summary requirement provides a workable model, and jurisdictions without such a rule should adopt one. A machine-readable registry of rights reservations would reduce the current uncertainty over whether robots.txt instructions or website terms count as valid opt-outs.

The second is statutory remuneration. Individual licensing across billions of works is impractical, and outright prohibition would concentrate AI development in whichever countries permit training. A licence by law, under which developers may train on published works in return for a royalty distributed through collective management organisations, would give developers certainty and creators income. Senfleben (2023) makes a similar case for remuneration at the output stage. The approach has precedents in statutory licensing for broadcasting and private copying levies. It would not cover memorised reproduction of the kind found in *GEMA*, which should remain infringing, because a model that can regenerate a work supplies a substitute for it.

The third is a limited right in AI-generated output. Denying all protection creates a free-rider problem, but granting full copyright to machine output for the life of an author plus seventy years would be excessive. A neighbouring right lasting around fifteen years, similar to the sui generis database right under Directive 96/9/EC, vested in the person who makes the arrangements for generation and without moral rights, would protect investment without pretending that a machine is an author. The United Kingdom could achieve this by shortening the term in section 12(7), and other jurisdictions could legislate for it directly.

None of these reforms will work well if adopted piecemeal. WIPO's Standing Committee on Copyright and Related Rights, which has discussed AI since 2019, is the natural forum for agreeing minimum standards on transparency and remuneration (WIPO, 2020). A soft-law instrument that states implement domestically would be

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

more realistic than a new treaty. Objections can be expected from both sides. Some creators will say that a statutory licence removes their right to refuse; developers will say that any payment slows innovation. The first objection has most force for moral rights and personal likeness, which should remain outside the licence. The second is weakened by the size of the settlement in *Bartz*, which shows that the cost of legal uncertainty is already very high.

Conclusion

Generative AI has not made intellectual property law obsolete, but it has exposed assumptions that were never tested. Copyright assumed that every protected work had a human author and that copying was an act of exploitation that could be seen and priced. Patent law assumed that invention happens in a human mind. The cases decided between 2023 and 2026 show that courts can apply these assumptions consistently to the clearest cases. They are less able to deal with the middle ground, where people use AI systems as collaborators, or with the economic question of who should benefit from models built on the work of others.

The comparative evidence points in a consistent direction. Output rules are converging around human creative control, patent rules have converged on human inventorship, and training rules remain fragmented. That fragmentation is the most serious problem, because it affects the largest sums and is the easiest to exploit by moving activity across borders. Transparency, statutory remuneration and a short output right would not resolve every difficulty, but they would give creators a share of the value their work produces and give developers the certainty they need to invest. Future research should follow the appeals in *GEMA* and *Getty*, the pending American litigation and the implementation of the AI Act's transparency duties, since these will show whether the law can adjust through interpretation or whether legislatures will have to act.

Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors thank colleagues at Northwest University, the University of Lincoln and Fatima Jinnah Women University for comments on earlier drafts. Any remaining errors are the authors' own.

References

- Abbott, R. (2016). I think, therefore I invent: Creative computers and the future of patent law. *Boston College Law Review*, 57(4), 1079–1126.
- Abbott, R. (2020). *The reasonable robot: Artificial intelligence and the law*. Cambridge University Press.
- Agency for Cultural Affairs, Government of Japan. (2024). *General understanding on AI and copyright in Japan*. Legal Subcommittee under the Copyright Subdivision of the Cultural Council.
- Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201-WHO (N.D. Cal. filed Jan. 13,

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

2023).

Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith, 598 U.S. 508 (2023).

Authors Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015).

Bartz v. Anthropic PBC, No. 3:24-cv-05417-WHA (N.D. Cal. June 23, 2025) (order on fair use).

Bridy, A. (2012). Coding creativity: Copyright and the artificially intelligent author. *Stanford Technology Law Review*, 2012, 5.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53 (1884).

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

Carlini, N., Tramèr, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., et al. (2021). Extracting training data from large language models. In *Proceedings of the 30th USENIX Security Symposium* (pp. 2633–2650).

Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwag, V., Tramèr, F., et al. (2023). Extracting training data from diffusion models. In *Proceedings of the 32nd USENIX Security Symposium* (pp. 5253–5270).

Commissioner of Patents v. Thaler [2022] FCAFC 62.

Copyright, Designs and Patents Act 1988 (UK), c. 48.

Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market. *Official Journal of the European Union*, L 130, 92–125.

European Patent Office, Legal Board of Appeal. (2021). Decisions J 8/20 and J 9/20 (DABUS), 21 December 2021.

Feist Publications, Inc. v. Rural Telephone Service Co., 499 U.S. 340 (1991).

GEMA v. OpenAI, Landgericht München I [Regional Court Munich I], Case No. 42 O 14139/24 (Nov. 11, 2025).

Gervais, D. J. (2020). The machine as author. *Iowa Law Review*, 105(5), 2053–2106.

Getty Images (US) Inc. v. Stability AI Ltd. [2025] EWHC 2863 (Ch).

Ginsburg, J. C., & Budiardjo, L. A. (2019). Authors and machines. *Berkeley Technology Law Journal*, 34(2), 343–456.

Guadamuz, A. (2017). Do androids dream of electric copyright? Comparative analysis of originality in artificial intelligence generated works. *Intellectual Property Quarterly*, 2017(2), 169–186.

Henderson, P., Li, X., Jurafsky, D., Hashimoto, T., Lemley, M. A., & Liang, P. (2023). Foundation models and fair use. *Journal of Machine Learning Research*, 24(400), 1–79.

Hristov, K. (2017). Artificial intelligence and the copyright dilemma. *IDEA: The IP Law Review*, 57(3), 431–454.

Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base: Analyst note. *Reuters*.

Hutchinson, T., & Duncan, N. (2012). Defining and describing what we do: Doctrinal legal research. *Deakin Law Review*, 17(1), 83–119.

Infopaq International A/S v. Danske Dagblades Forening, Case C-5/08, EU:C:2009:465.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

- Kadrey v. Meta Platforms, Inc.*, No. 3:23-cv-03417-VC (N.D. Cal. June 25, 2025).
- Lee, K., Cooper, A. F., & Grimmelmann, J. (2023). Talkin' 'bout AI generation: Copyright and the generative-AI supply chain. *arXiv preprint arXiv:2309.08133*.
- Lemley, M. A., & Casey, B. (2021). Fair learning. *Texas Law Review*, 99(4), 743–785.
- Li v. Liu*, Beijing Internet Court, (2023) Jing 0491 Min Chu No. 11279 (Nov. 27, 2023).
- Naruto v. Slater*, 888 F.3d 418 (9th Cir. 2018).
- Nova Productions Ltd v. Mazooma Games Ltd* [2007] EWCA Civ 219.
- Painer v. Standard VerlagsGmbH*, Case C-145/10, EU:C:2011:798.
- Quintais, J. P. (2025). Generative AI, copyright and the AI Act. *Computer Law & Security Review*, 56, 106107.
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 12 July 2024.
- Sag, M. (2023). Copyright safety for generative AI. *Houston Law Review*, 61(2), 295–347.
- Samuelson, P. (2023). Generative AI meets copyright. *Science*, 381(6654), 158–161.
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., et al. (2022). LAION-5B: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35, 25278–25294.
- Senftleben, M. (2023). Generative AI and author remuneration. *IIC: International Review of Intellectual Property and Competition Law*, 54, 1535–1560.
- Shenzhen Tencent Computer System Co. v. Shanghai Yingxun Technology Co.*, Shenzhen Nanshan District People's Court, (2019) Yue 0305 Min Chu No. 14010 (Dec. 24, 2019).
- Thaler v. Comptroller-General of Patents, Designs and Trade Marks* [2023] UKSC 49.
- Thaler v. Perlmutter*, 687 F. Supp. 3d 140 (D.D.C. 2023), aff'd, 130 F.4th 1039 (D.C. Cir. 2025), cert. denied, No. 25-449 (U.S. Mar. 2, 2026).
- Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022), cert. denied, 143 S. Ct. 1783 (2023).
- The New York Times Co. v. Microsoft Corp.*, No. 1:23-cv-11195 (S.D.N.Y. filed Dec. 27, 2023).
- Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc.*, No. 1:20-cv-00613-SB (D. Del. Feb. 11, 2025).
- UK Intellectual Property Office. (2024). *Copyright and artificial intelligence: Consultation* (published 17 December 2024). UK Government.
- U.S. Copyright Office. (2023a). *Re: Zarya of the Dawn* (Registration No. VAu001480196), letter of 21 February 2023.
- U.S. Copyright Office. (2023b). Copyright registration guidance: Works containing material generated by artificial intelligence. *Federal Register*, 88(51), 16190–16194.
- U.S. Copyright Office. (2025a). *Copyright and artificial intelligence, Part 2: Copyrightability*. Library of Congress.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

- U.S. Copyright Office. (2025b). *Copyright and artificial intelligence, Part 3: Generative AI training* (pre-publication version). Library of Congress.
- U.S. Patent and Trademark Office. (2025). *Revised inventorship guidance for AI-assisted inventions* (rescinding the guidance of 13 February 2024). November 2025.
- World Intellectual Property Organization. (2020). *Revised issues paper on intellectual property policy and artificial intelligence* (WIPO/IP/AI/2/GE/20/1 REV). WIPO.
- Yanisky-Ravid, S. (2017). Generating Rembrandt: Artificial intelligence, copyright, and accountability in the 3A era. *Michigan State Law Review*, 2017(4), 659–726.
- Zweigert, K., & Kötz, H. (1998). *An introduction to comparative law* (3rd ed., T. Weir, Trans.). Oxford University Press.