(2025) 13 NIBLeJ 9

Getty Images v Stability AI [2025] EWHC (Ch): A Case Note on Generative AI, Copyright, and the Imperative of Global Collaboration

Xiaocong LIU,[1] Kai ZHANG[2] and Rongtao ZHANG[3]

Abstract

This case note critically analyses the High Court's landmark judgment in Getty Images (US) Inc v Stability AI Ltd [2025] EWHC (Ch). The ruling reveals a stark conflict between territorial copyright regimes and the borderless operational mechanisms of generative artificial intelligence. Whilst the court dismissed the plaintiff's principal claims of copyright and database rights infringement, its ruling was primarily grounded in jurisdictional limitations. Specifically, the identification of the model's training location could not absolve liability for the training process itself. Furthermore, the judiciary's refusal to characterise the trained model as an 'infringing article' established a crucial distinction between the transient copying inherent in machine learning and the permanent reproduction prohibited under the Copyright, Designs and Patents Act 1988. By dissecting the court's reasoning, this paper argues that the judgment vividly exposes the 'regulatory arbitrage' phenomenon spawned by today's fragmented legal landscape. Consequently, the author contends that domestic reforms are wholly inadequate, urging urgent global collaboration to harmonise legal standards, technical protocols, and institutional oversight. This is essential to prevent the governance of transnational AI ecosystems from descending into a vicious cycle of 'race to the bottom' in regulatory standards.

Key words

Generative AI; Copyright Infringement; Territorial Jurisdiction; CDPA 1988; Intellectual Property Law.

[1] LLM, PhD, Nottingham Law School, Nottingham Trent University.

[2] Law School and Intellectual Property School, Jinan University. The author can be contacted via kaizhanglaw@gmail.com.

[3] Bristol University.

Introduction and Case Overview

On 4 November 2025, the High Court of Justice, Business and Property Courts of England and Wales, Intellectual Property List (ChD), handed down its long-anticipated judgment in *Getty Images (US) Inc and others v Stability AI Ltd*([2025] EWHC (Ch)),[4] delivered by Mrs Justice Joanna Smith DBE. The decision marked the first comprehensive judicial engagement in the United Kingdom with the legal ramifications of generative artificial intelligence (AI), particularly the question of whether the training and deployment of image-generation models constitutes copyright infringement, database right infringement, or trademark infringement under existing law. The case, brought by a consortium of Getty Images entities incorporated in the United States, Ireland, the United Kingdom and Canada, epitomised the transnational character of AI development and the inadequacy of territorially bounded intellectual property frameworks in addressing such globally distributed technologies.

The claimants alleged that Stability AI, the developer of the image-generation model *Stable Diffusion*, had unlawfully used millions of Getty-owned or licensed images to train its generative AI system. Getty claimed that Stability's scraping and reproduction of these images, and the storage of copies for the purpose of training, constituted primary copyright infringement under sections 16 and 17 of the Copyright, Designs and Patents Act 1988 (CDPA). It further alleged that the trained model itself, when imported and distributed in the United Kingdom, constituted an "infringing article" within the meaning of sections 22 and 23 CDPA, giving rise to secondary infringement. Additional claims were made for infringement of database rights under the Copyright and Rights in Databases Regulations 1997, for trademark infringement under the Trade Marks Act 1994, and for passing off. The plaintiffs pointed to outputs generated by *Stable Diffusion* that included faint or partial reproductions of Getty's distinctive watermark ("Getty Images" or "iStock") as evidence of both copyright copying and trademark confusion.

Stability AI denied the allegations. It contended that *Stable Diffusion* does not store or reproduce copyrighted images in any recoverable form but instead learns high-level statistical correlations between visual features and textual descriptions. It further argued that the training process took place entirely outside the United Kingdom, using servers located in the United States and elsewhere, and therefore fell outside the territorial jurisdiction of UK copyright law. Moreover, the defendant maintained that any apparent

[4] 1. Getty Images (US) Inc and others v Stability AI Ltd [2025] EWHC (Ch) (High Court of Justice, Business and Property Courts of England and Wales, Intellectual Property List (ChD), Mrs Justice Joanna Smith DBE, 4 November 2025).

reproductions of watermarks were rare, unintentional artefacts produced through overfitting, not the result of deliberate copying or use of Getty's trademarks in trade.[5]

After extensive procedural delays, expert evidence, and cross-examination, Mrs Justice Smith issued judgment largely in favour of Stability AI. The court dismissed Getty's claims of copyright and database right infringement, rejected the argument that the *Stable Diffusion* model was itself an infringing article, and found only a narrow instance of trademark infringement, insufficient to support the broader claim of passing off. The judgment is therefore both a doctrinal milestone and a regulatory wake-up call: while affirming that existing intellectual property law struggles to accommodate the mechanics of generative AI, it simultaneously exposes the structural limits of national regulation for a technology whose operations transcend borders.

This case note analyses the decision in three dimensions: first, the reasoning and key findings of the court; second, the implications for rights-holders, developers, and policymakers; and third, the normative and structural lesson that emerges, namely, the pressing need for global collaboration in the governance of generative AI.

The Court's Reasoning and Its Implications

The judgment turned on two central issues: first, whether training a generative AI model on copyrighted material constitutes copying within the meaning of the CDPA; and second, whether the trained model or its outputs can be treated as infringing articles or reproductions of those works. The court's conclusions on these points have far-reaching consequences for both copyright doctrine and the broader regulatory landscape of AI.

The claim for primary infringement failed at the threshold of territorial jurisdiction. Getty conceded that it could not prove that any act of copying or downloading of its images occurred within the United Kingdom. Evidence indicated that the scraping and training processes took place on remote servers located outside the jurisdiction, particularly in the United States.[6] Under orthodox principles of UK copyright law, infringement must occur within the territory to fall within the jurisdiction of the court. Because the claimants could not establish that any acts of reproduction under section 17 CDPA had occurred in the UK, the court declined to consider whether the training

[5] Ropes & Gray LLP, *Getty Image Loses Copyright Infringement Claim Against Stability AI in UK's First Generative AI Decision* (4 November 2025)
[6] Pinsent Masons LLP, *Getty's Copyright Case Against Stability AI Fails as High Court Rules Training Occurred Outside UK Jurisdiction* (4 November 2025)

process would, in principle, constitute infringement had it been performed domestically. This procedural outcome leaves the substantive question unresolved: does the ingestion of copyrighted works into the training data of a machine-learning model amount to reproduction for the purposes of copyright law? The High Court deliberately left this point open, although the reasoning of Mrs Justice Smith hints that, even had jurisdiction been established, the nature of model training may have complicated the infringement analysis.

The claim for secondary infringement, importation or possession of an "infringing article", was decided on substantive grounds. Getty argued that the trained model itself embodied infringing copies of its photographs because it was derived from, and mathematically encoded, those images. The court rejected this contention, drawing on expert testimony that the model's internal parameters, or "weights," do not store pixel-level data from training images but represent abstract numerical relationships extracted through statistical optimisation.[7] Accordingly, while Getty's images were undoubtedly copied during the training process, the resulting model did not constitute a copy or reproduction of those works. The model was a new, intangible object generated by a process of learning, not an infringing derivative. Thus, importing or using the model in the UK did not infringe copyright or database rights under sections 22 or 296 of the CDPA.

This reasoning effectively distinguishes between the *ephemeral copying* inherent in machine learning and the *permanent reproduction* that copyright law traditionally prohibits. It acknowledges that while training involves transient storage and analysis of works, those acts occur as part of a computational process whose outputs are probabilistic models rather than fixed reproductions. The court's willingness to conceptualise "copying" in this nuanced way marks a significant evolution in copyright jurisprudence. It suggests that the "copy" in machine-learning contexts cannot be defined merely by technical duplication but must consider the functional transformation of data into abstract model parameters.[8]

The court, however, did find limited trademark infringement under the Trade Marks Act 1994. Getty provided evidence that certain images generated by *Stable Diffusion* bore visual artefacts resembling its watermark. Mrs Justice Smith accepted

[7] Osborne Clarke LLP, *Stability AI Generates a Big Win in Getty Images UK Copyright Action* (5 November 2025) https://www.osborneclarke.com/insights/getty-v-stability-ai-stability-ai-generates-big-win-english-courts-landmark-first-judgment

[8] William Fry Solicitors, *Getty Images v Stability AI: The Most Important AI Legal Decision to Date* (6 November 2025) https://www.williamfry.com/knowledge/getty-images-v-stability-ai-the-most-important-ai-legal-decision-to-date/

that the appearance of a distinctive mark capable of signifying origin could, in principle, constitute use in the course of trade if disseminated to consumers. Yet she found the evidence of actual use too limited to support a broad injunction. The watermark reproductions occurred primarily in contrived tests rather than in publicly distributed outputs, and there was no evidence of consumer confusion or economic harm.[9] Consequently, while Stability AI was technically liable for limited trademark infringement, the court declined to grant substantial remedies.

From a doctrinal perspective, the judgment establishes several important propositions. First, it confirms that existing copyright concepts struggle to map onto machine-learning processes, which involve abstraction and transformation rather than duplication. Second, it reinforces the territoriality of copyright law: even global digital processes remain governed by the place where the act of copying occurs. Third, it illustrates that trademark law may still operate as a residual check on generative AI outputs that misappropriate protected identifiers, though its application remains narrow. Finally, it reveals that the present legal framework provides little clarity or recourse for rights-holders whose works are used to train models outside their home jurisdictions.

The implications of these findings are profound. For AI developers, the judgment offers temporary legal certainty. Under current UK law, training a model on copyrighted material abroad does not automatically expose the resulting model to domestic liability, so long as it does not contain retrievable reproductions or generate infringing outputs. For rights-holders, by contrast, the decision represents a setback: it demonstrates that traditional copyright law offers limited protection against unlicensed use of creative works for AI training, particularly when the acts occur outside their jurisdiction. For policymakers, the case signals that domestic legal systems may be ill-equipped to handle technologies that are inherently transnational and probabilistic. The outcome, therefore, both relieves and alarms, it gives developers breathing space while exposing the inadequacy of the legal architecture that governs them.

The Need for Global Collaboration in Regulating Generative AI

The *Getty Images v Stability AI* judgment ultimately does more than interpret existing law; it exposes the structural deficiencies of national legal systems in addressing a

[9] Mayer Brown LLP, *Getty Images v Stability AI: What the High Court's Decision Means for Rights-Holders and AI Developers* (7 November 2025) https://www.mayerbrown.com/en/insights/publications/2025/11/getty-images-v-stability-ai-what-the-high-courts-decision-means-for-rights-holders-and-ai-developers

global technological phenomenon. The decision illustrates how jurisdictional limits, doctrinal rigidity, and institutional fragmentation combine to produce legal uncertainty and uneven accountability. This final section argues that meaningful governance of generative AI cannot be achieved through isolated national measures, but instead requires coordinated global collaboration across legal, technical, and institutional domains.

Generative AI development is inherently transnational. Datasets are scraped from the global internet, stored on distributed cloud servers, and processed using compute clusters that may be located in multiple jurisdictions. Model outputs are disseminated globally through open-source platforms and APIs. In such an ecosystem, the territorial scope of national copyright or trademark law, designed for a world of localized acts and physical media, becomes conceptually obsolete. The *Getty* case exemplifies this disjuncture: the decisive issue was not whether training constituted copying, but where it occurred. The fact that acts of copying transpired outside the UK placed them beyond the court's reach, even though the effects, in the form of a globally distributed model, were felt within it. As Mrs Justice Smith observed, the territorial logic of copyright law sits uneasily with the distributed architecture of AI training.[10]

This jurisdictional fragmentation invites what scholars term "regulatory arbitrage": developers can avoid liability by situating training operations in lenient jurisdictions, while still profiting from global deployment. Absent international coordination, such arbitrage risks precipitating a "race to the bottom," as states compete to attract AI investment by offering permissive environments. The result is a patchwork of inconsistent rules, uncertain rights, and weakened incentives for creators. This dynamic mirrors earlier stages of the digital economy, where inconsistent data protection and tax regimes enabled transnational avoidance until harmonization efforts such as the EU's General Data Protection Regulation set global benchmarks.

Global collaboration in regulating generative AI must therefore rest on three pillars: harmonization of legal standards, coordination of technical practices, and institutional mechanisms for accountability. The first pillar, legal harmonization, requires states to negotiate common principles governing the use of copyrighted works in AI training. Existing international instruments, such as the Berne Convention, the WIPO Copyright Treaty, and the TRIPS Agreement, already provide a framework for mutual recognition of rights. However, these instruments predate machine learning and make no provision

[10] Lexology, *High Court Rejects Getty Images' Copyright Claims Against Stability AI and Clarifies Jurisdictional Limits* (10 November 2025) https://www.lexology.com/library/detail.aspx?g=97ec6470-9c7b-4249-a61d-fea9752882f0

for data-driven abstraction. A new international instrument, whether a protocol to Berne or a standalone treaty, could specify the conditions under which copyrighted materials may be used for AI training, balancing the rights of creators with the public interest in innovation. It might, for example, recognise a limited exception for text and data mining subject to opt-out or compulsory licensing mechanisms, akin to the approach taken in Articles 3 and 4 of the EU Directive on Copyright in the Digital Single Market (Directive (EU) 2019/790). Such harmonization would prevent developers from exploiting gaps between jurisdictions and would give creators a predictable framework for compensation.

The second pillar concerns technical coordination. Legal rules alone cannot ensure compliance unless they are translated into technical design choices. International standards bodies, industry associations, and research consortia should therefore develop shared technical protocols for transparency, data provenance, and content attribution. These could include obligations for AI developers to disclose the datasets used for training, to implement watermarking or metadata tags in generated content, and to employ techniques that reduce the risk of memorizing or reproducing copyrighted works. Such transparency would enable rights-holders to identify when their works have been used and to exercise opt-out or compensation rights. The *Getty* judgment implicitly underscores this need: had Stability AI been required to disclose its data sources or training methods in a standardized format, many of the evidentiary uncertainties in the litigation could have been avoided. Technical standards thus complement legal rules by embedding accountability into the architecture of AI systems.

The third pillar is institutional cooperation. Because AI development and deployment transcend borders, enforcement cannot depend solely on national courts. International or transnational bodies — possibly under the auspices of the World Intellectual Property Organization (WIPO), the Organisation for Economic Co-operation and Development (OECD), or a newly constituted Global AI Council — could be empowered to oversee compliance, mediate disputes, and coordinate enforcement across jurisdictions. These institutions could maintain registries of licensed datasets, monitor adherence to transparency requirements, and facilitate cross-border remedies. Without such coordination, rights-holders like Getty will face insurmountable hurdles in pursuing claims across multiple legal systems, each with its own evidentiary standards and procedural rules.

Global collaboration must also extend beyond copyright to encompass broader concerns of accountability and informed governance. Generative AI implicates not only

intellectual property but also misinformation, deepfakes, data protection, and cultural diversity. Regulation must therefore adopt a holistic, multi-stakeholder approach that integrates human rights, ethics, and competition policy. The *Getty* decision demonstrates that focusing narrowly on copyright misses the systemic nature of the challenge. A model trained on billions of internet images does not merely raise issues of authorship or ownership; it reshapes the informational ecosystem by automating the creation and circulation of visual culture. Effective governance must therefore address the epistemic and democratic dimensions of AI, ensuring that the systems which produce and disseminate knowledge are subject to transparency and oversight.

From a normative standpoint, global collaboration should aim to operationalise the principles of *accountable AI* and *informed governance*. Accountable AI entails that those who design, deploy, and profit from AI systems bear responsibility for their social and legal impacts. This responsibility requires traceability of training data, transparency of model behaviour, and mechanisms for redress when harm occurs. Informed governance, by contrast, denotes the capacity of regulatory institutions and democratic publics to understand, evaluate, and steer AI systems in the public interest. Open-source development, transparency mandates, and international cooperation all serve this goal by reducing information asymmetries between private developers and public regulators. In this sense, open-source AI can function as a civic infrastructure: by making models auditable and data flows visible, it strengthens the epistemic foundations of democratic oversight in the digital age. Such openness, however, must be coupled with global norms to prevent the exploitation of open resources by actors who operate beyond the reach of national law.

The High Court's decision thus crystallizes a paradox. On one hand, it upholds the integrity of existing legal doctrine by refusing to stretch copyright concepts beyond their logical limits. On the other, it exposes how those very limits leave rights-holders unprotected and society ill-prepared for the scale of generative AI. The judgment's restraint is commendable in legal method but troubling in policy consequence. It is now incumbent on lawmakers, regulators, and international organizations to fill the vacuum that courts cannot. Domestic reforms (e.g. clarifying the status of text and data mining, establishing licensing mechanisms for AI training, and mandating transparency for dataset composition) are necessary first steps. Yet without international coordination, such reforms risk being undercut by extraterritorial operations and divergent national practices.

The future of generative AI regulation must therefore be global by design. An effective framework would combine harmonized legal norms, interoperable technical standards, and cooperative institutions. It would recognise the legitimate interests of creators while preserving the public benefits of AI innovation. It would distribute responsibility among developers, users, and intermediaries, ensuring that accountability does not end at the borders of any single state. And it would embed transparency as a precondition of trust, allowing societies to govern AI not through fear or prohibition but through informed participation.

The *Getty Images v Stability AI* decision is thus not the final word but the beginning of a larger conversation. It reveals that the law, operating within national confines, cannot alone reconcile the competing imperatives of creativity, innovation, and fairness in a borderless technological ecosystem. The case should be read as a call to collective action, to build a global regulatory order that matches the global scale of the technology it seeks to govern. Without such collaboration, the world risks entrenching a fragmented regime where generative AI operates in legal limbo, rights-holders are left uncompensated, and accountability dissolves across borders. With collaboration, however, the promise of generative AI, as a tool for human creativity, knowledge, and shared progress, can be realised within a framework of legitimacy and justice.