
# Much ado about little

The *Getty Images* judgment fails to deliver on expectations. Experts from the IP & Technology teams at Arnold & Porter explain why

© Algi Febri Sugita/ZUMA Press Wire/Shutterstock

**IN BRIEF**

- Summarises the highlights of the court's findings in *Getty Images v Stability AI*— a 'historic' but 'extremely limited' ruling.
- Sets out the impact the judgment will have on future AI infringement cases.

A judge in the High Court of Justice is the latest to weigh in on liability for use of intellectual property by AI companies, in *Getty Images (US) Inc and others v Stability AI Ltd* [2025] EWHC 2863 (Ch). Mrs Justice Joanna Smith DBE considered claims that (i) Stability AI infringed Getty Images' registered trade marks by generating synthetic image outputs containing similar marks using its AI model, Stable Diffusion, and (ii) Stability AI infringed the copyright in Getty's visual content by using Getty's images to train models of Stable Diffusion.

In a painstakingly detailed 205-page judgment, Mrs Justice Joanna Smith made what she called a 'historic' but 'extremely limited' ruling that a handful of synthetic image outputs from Stable Diffusion infringed Getty Images' trade mark rights under two sections of the Trade Marks Act 1994 (TMA 1994). The court dismissed Getty Images' remaining claims of trade mark dilution, tarnishment or unfair advantage; passing off; and secondary copyright infringement.

## Trade mark claim findings & conclusions of law

Getty's trade mark claims arose from allegations that Stability Diffusion generated images bearing watermarks that infringed Getty's registered trade marks, 'Getty Images' and 'iStock.' Getty presented examples of synthetic images generated by Stable Diffusion that featured watermarks similar to Getty's marks. The appearance of these watermarks was caused by Stability AI's unlicensed use of images owned or exclusively licensed by Getty to train Stable Diffusion, causing the model to 'memorize' Getty's marks and generate the watermarks in its output.

The court considered Getty's infringement claims under three sections of TMA 1994:

(i) s 10(1), which requires demonstrating 'double identity', meaning that the an infringement occurs when an allegedly infringing mark is identical to the trade mark, and the infringing mark is used to market an identical good or service;
(ii) s 10(2), which requires demonstrating that the marks are similar and likely to create consumer confusion; and
(iii) s 10(3), which requires demonstrating similarity between the marks plus mark dilution, tarnishment, or unfair advantage to the infringer.

Getty originally asserted an additional claim for passing off, but the parties did not fully brief the claim and the court declined to decide it.

The court concluded that one of the images submitted to the court infringed the 'iStock' marks under s 10(1), and that three of the images infringed the 'Getty Images' mark under s 10(2). The court found no violation of s 10(3), and no evidence of violations on a 'widespread scale' (at [756]). The judgment did not state the amount of damages the court would award to Getty as UK cases consider liability in an initial trial before consideration of financial relief in a separate inquiry as to damages or account of profits trial.

## Real-world users

The court spent a significant portion of the ruling deciding the threshold issue of whether Getty satisfied its burden of demonstrating that the watermarks were presented to real-world users in the UK, as required under TMA 1994. The parties submitted three types of evidence on this issue: a random sample of 1,000 images selected by Stability's experts, which were generated by real-world users and showed no instances of the watermarks; thousands of experimentally created examples generated by Getty's experts, all of which featured the watermarks; and 26 examples of real-world instances of images with the watermarks, mostly from social media websites such as Reddit, with no evidence of the user's geographic location.

Stability argued its sample of 1,000 images demonstrated it was statistically unlikely that a user in the UK would be presented with the watermark. The court largely rejected the sample as evidence, agreeing with testimony from experts that the sample size was too small and that a sample size of about a million would make it a reliable indicator of the probability the model would output an infringing image (at [198]). Moreover, the sample images were 'of little assistance on the question of whether real world users in the UK have in fact generated watermarks', particularly in light of the real-world examples of images with the marks and Getty's experimentally induced images (at [199]).

The court also rejected most of Getty's experimentally produced images, which were generated by prompts consisting of verbatim descriptions of images that were copied and pasted from Getty's website. Because Getty presented no proof that real-world users would use such a prompt, the court agreed with Stability that the images were too 'contrived' to count as evidence of any real-world incidence of images with the watermarks. On the other hand, the court accepted Getty's experimentally produced images that were generated using prompts containing the phrases 'news photo' and 'vector art'. The court found there was sufficient evidence that these prompts were 'representative of prompts that real world users would use', and thus the images were evidence that similar images were likely to be generated by users in the UK (at [218]).

Finally, Getty submitted an exhibit of

26 allegedly infringing images that were generated by real-world users using Stable Diffusion, but admitted that there was no evidence that any of the images were generated by users within the UK. However, Getty argued that 'evidence that it has happened in the real world somewhere is equally good evidence that it is likely to have happened in the UK' (at [157]). The court agreed, finding the real-world images were persuasive to show that similar images were likely to be generated by users inside the UK.

In sum, while Getty could not rely on thousands of 'contrived' experimental images as the basis of its infringement claim, the court concluded it was likely that UK users had been presented with infringing images based on (i) representative images generated by experimental prompts that were similar to prompts likely to be used by real-world users in the UK; and (ii) examples of real-world instances of images with the watermarks, even where those examples were not generated in the UK.

## Identity or similarity

The court considered the aural, visual and conceptual similarities and differences between the watermarks in Stability's synthetic images and Getty's registered marks. The court ultimately found only one of the sample images—the Spaceship Image shown above—contained watermarks identical to one of Getty's registered marks.

As to the Japanese Temple Garden Image, shown above, the court rejected Getty's submission that the watermark was identical to Getty's registered mark because it had 'an extra 'i' in the word 'images' so that it appears to read 'imaiges' (at [414]). The court found that this difference was not 'so insignificant that it may go unnoticed by an average consumer exercising a moderate degree of attention' (at [414]). However, for purposes of similarity under s 10(2), the court found the Japanese Temple Garden Image was 'highly similar' to Getty's 'Getty Images' mark, finding similarity was satisfied despite the presence of the extra 'i' (at [459]).

## Likelihood of confusion

In the court's s 10(2) analysis of likelihood of consumer confusion, it found that the average consumer was likely to believe that synthetic images with watermarks similar to Getty's marks were associated with Getty. This was the case both for unsophisticated users, whom the court found would likely believe Getty itself provided the generated image, and more 'tech-savvy' users, whom the court found would likely believe Stability licensed Getty's images or at least obtained permission to use them to train Stable Diffusion.

## Dilution, tarnishment & unfair advantage

The court rejected Getty's claim under the remaining portion of TMA 1994, s 10(3), which required showing similarity of the marks in addition to dilution, tarnishment or unfair advantage. The court's conclusion was based on evidence that Stable Diffusion consumers do not desire watermarks on synthetic images and that Stability actively took steps to prevent the generation of images with watermarks. The court also reasoned that it would not make sense for a consumer who wished to avoid paying a licensing fee to use a Getty image to instead use an image generated by Stable Diffusion that bears a watermark.

Getty also argued that its mark was damaged by the generation of 'not safe for work' or pornographic images. The court was more open to this theory of harm, but ultimately concluded that there was insufficient evidence that any 'not safe for work image' was generated by a user within the UK.

> *"The evidentiary requirement for showing intellectual property infringement in AI outputs can be difficult to meet"*

## Secondary copyright & database rights

Getty Images had originally brought a number of copyright claims against Stability AI, including:

(i) a 'training and development claim' alleging that Stability AI scraped millions of Getty Images' copyrighted visual assets and used them to develop Stable Diffusion;
(ii) an 'outputs claim' alleging that the images produced by Stable Diffusion infringed Getty's copyright; and
(iii) a 'database rights infringement claim' alleging database rights infringement.

Because (i) all of the training took place outside of the UK and (ii) Stability AI updated Stable Diffusion so that it now blocks the prompts used to generate the alleging infringing output, Getty acknowledged that all three of these claims could no longer be advanced (at [9]).

Given the above, only one of Getty's two original secondary infringement claims under ss 22 and 23 of the Copyright, Designs and Patents Act 1988 (CDPA 1988) remained for the court to decide: the claim that the Stable Diffusion model is an 'infringing' copy because it was imported into the UK and, had it been created in the UK, would have constituted infringement under s 27(3) of CDPA 1988.

There are two issues of law related to this secondary infringement claim: whether Stable Diffusion is capable of being (i) an 'article' under ss 22 and 23 of CDPA 1988; and (ii) an 'infringing copy' under s 27(3) of CDPA 1988. Getty argued that the definition of 'infringing copy' in s 27(3) was broad enough to include an intangible article such as Stable Diffusion and that Stable Diffusion's model weights are the 'article[s]', and that therefore 'making' the model weights would have constituted infringement had it been done in the UK.

Stability argued that 'article' is limited to tangible objects and that 'infringing copy' cannot apply to Stable Diffusion to the extent the model was trained in the US and did not include UK copyrighted works (at [550]).

The court interpreted the word 'article' to be capable of being 'an electronic copy stored in an intangible medium (such as the AWS Cloud)' because it is 'in [the court's] judgment, capable of being an infringing copy' (at [583]) and [590]. In other words, an 'article' cannot be divorced from the concept of an infringing copy (at [570]–[590]). The court relied on the statutory construction of s 17 of CDPA 1988 (which describes 'Infringement of copyright by copying') to give context to the interpretation of ss 22, 23 and 27(3) and found persuasive the interpretations of 'article[s]' in other cases, where they encompassed electronic storage media, including, for example, that a random access memory (RAM) chip was capable of being an article and that copies of a 'work in storage space made available to a user in connection with cloud computing services constitutes a reproduction of that work' (at [579] and [582]).

Next, the court analysed whether the Stable Diffusion model constitutes 'an infringing copy' (at [592]). The court recalled that it was 'common ground' that 'the model weights of the various Stable Diffusion versions do not store the visual information in the Copyright Works',

that is, there is no present copy of the copyrighted works in the model weights.

Stability argued that, because no copy was stored in the model weights, the Stability Diffusion models do not amount to an infringing copy. By contrast, Getty argued that 'there is no requirement that the article thus made must continue to retain a copy or copies of the work'; instead, 'copying' included 'the making of copies which are transient or are incidental to some other use of the work' (at [593]; emphasis in original judgment). In other words, Getty's argument was that, as soon as it was made, the AI model was an 'infringing copy' (at [594]).

Relying on *G4S Plc v G4S Trustees Ltd* [2018] EWHC 1749 (Ch), [2018] All ER (D) 75 (Jul) and *Sony Computer Entertainment Inc v Ball* [2004] EWHC 1738 (Ch), [2005] FSR 9, the court held that 'an infringing copy must be a copy… I cannot see how an article can be an infringing copy if it has never consisted of/stored/contained a copy… [A]n article becomes an infringing copy when the act of reproduction occurs. From that moment the article is an infringing copy—but it ceases to be an infringing copy once it no longer contains the copy' (at [597]).

The court summarised that the parties' dispute turned on whether an article whose creation involved the use of infringing copies, but which never contained or stored those copies, was itself an infringing copy such that its making in the UK would have constituted an infringement. The court concluded that it was 'not enough… that… "the time of making of the copies of the Copyright Works **coincides** with the making of the Model' (emphasis added)… While it is true that the model weights are altered during training by exposure to Copyright Works, by the end of that process the Model itself does not store any of those Copyright Works; the model weights are not themselves an infringing copy and they do not store an infringing copy.' The court thus dismissed Getty's claim of secondary infringement.

## Key takeaways

Probably the most significant takeaway from the judgment is the finding that, under English law, AI models are not, in and of themselves, infringing copies of their training data. This is largely consistent with the few court decisions in the US that have considered this issue. See:

- *Thomson Reuters Enterprise Centre GmbH v Ross Intelligence Inc*, 765 F.Supp.3d 382, 398 (D Del 2025): copies of copyrighted Westlaw headnotes used to train defendant's AI model were not 'part of the final product that [the defendant] put forward to customers'; rather, '[t]he copying occurred at an intermediate step: [the defendant] turned the headnotes into numerical data about the relationships among legal words to feed into its AI.').
- *Bartz v Anthropic PBC*, 787 F.Supp.3d 1007, 1022 (ND Cal 2025): training large language models by 'making copies' of copyrighted works 'within the LLM or otherwise' was non-infringing fair use.
- *Kadrey v Meta Platforms*, 788 F.Supp.3d 1026, 1045 (ND Cal 2025): holding copies of books made for purpose of training large language model was fair use where 'adversarial' prompts demonstrated that the model would output, at most, no 'more than 50 words of any of the plaintiffs' books'.

But also see *Andersen v Stability AI Ltd*, 700 F.Supp.3d 853, 864 (ND Cal. 2023): permitting plaintiff's copyright infringement claim against Stability AI to survive a motion to dismiss on the theory that Stability AI caused the images used for training 'to be stored at and incorporated into Stable Diffusion as compressed copies'.

However, another key takeaway is that the evidentiary requirement for showing intellectual property infringement in AI outputs can be difficult to meet: it requires obtaining strong real-world evidence of trade mark and copyright violations, not just theoretical possibilities or contrived examples. Proof that infringing images can possibly be generated with an AI model is insufficient to form the basis of a claim of trade mark infringement. Instead, complaining parties will need to tie infringing outputs to the actual behaviour of real-world users. Moreover, claims based on a rate of infringement may require large sample sizes, in the millions, which would entail significant costs in terms of computing power.

Finally, trade mark holders will need to show that their mark was infringed within the correct jurisdiction where the claim is brought. Getty largely failed to offer such jurisdiction-specific evidence, which resulted in Getty's inability to proceed with key copyright and database right infringement claims. That has not been an issue in the US, primarily because many of the leading AI developers have trained and offer their AI models within the US. **NLJ**

On 16 December 2025, the High Court granted Getty permission to appeal the dismissal of its secondary copyright infringement claim, while refusing Stability AI permission to appeal the trade mark infringement findings on the basis that any such appeal would have no real prospect of success. Stability AI may nevertheless seek permission to appeal directly from the Court of Appeal.

**Tom Magnani**, Technology Transactions practice head and Technology, Media & Telecommunications industry group co-chair, **Dr Beatriz San Martin**, partner, **Jessica Gillotte**, senior associate, & **Ally Myers,** associate, at Arnold & Porter (*arnoldporter.com*). *Newlawjournal.co.uk*