

Title: Getty Images v Stability AI: Unpacking the High Court's judgment

Author(s): [Nina O'Sullivan, Client Engagement Partner, Knowledge Lawyer](#)
nina.osullivan@mishcon.com

[Gareth Dickson, Partner](#)
gareth.dickson@mishcon.com

[Mark Dean, Associate](#)
mark.dean@mishcon.com

Full Article

In brief

- The High Court has rejected Getty Images' secondary copyright infringement claim against Stability AI, finding that Stable Diffusion's model weights - though trained on datasets that included Getty's images - did not constitute 'infringing copies' as they do not store the original copyright works. However, the court did find limited and historic trade mark infringement relating to watermarks appearing in outputs from certain versions of Stable Diffusion.
- Getty abandoned two significant claims during the trial: copyright infringement during the training and development process and copyright infringement in outputs. This means certain important issues relating to GenAI and copyright remain yet to be considered by UK courts. Infringement assessments will continue to be highly fact-sensitive.
- Balancing AI and creative industry interests remains a policy priority for the Government. Attention now turns to its response to its consultation on copyright and GenAI, as it faces pressure from creative industries opposing a general text and data mining exception that would allow AI companies to scrape copyright works unless rights holders expressly opt out.

Introduction

The High Court has delivered its judgment in [Getty Images v Stability AI](#), one of the most anticipated cases in recent years. Whilst the impact of the case had lessened due to Getty abandoning certain issues up to and during the trial, and Stability abandoning aspects of its defence, the court's decision considers a number of significant issues relating to the impact of GenAI tools on protection and enforcement of IP rights.

In summary, the court rejected Getty's claim of secondary copyright infringement. Stability AI's Stable Diffusion model was trained overseas and its importation into, and distribution within, the UK did not amount to dealing in an article which was an 'infringing copy'. In Getty's favour, however, the court did find trade mark infringement in relation to the display of watermarks in outputs from using certain iterations of Stable Diffusion, but this was "*historic*" and "*extremely limited*" in scope.

Background

Getty's business involves licensing visual assets for which it is either the owner or the exclusive licensee. At the heart of its case was the argument that UK company Stability had, from 2022 onwards, scraped millions of Getty's visual assets from its websites without its consent, and had used those images to train and develop various iterations of its Stable Diffusion GenAI model. The different iterations of Stable Diffusion have been made accessible via various means, including via the DreamStudio platform.

It was common ground that the training of Stable Diffusion involved downloading and storing copies of images obtained from relevant URLs on the AWS Cluster, which took place outside of the UK. Stability accepted that "at least some" images from Getty's websites were used during the training of Stable Diffusion.

To understand the issues in the case, it is important to have some appreciation of how Stable Diffusion works. It is a diffusion generative AI model, which involves model weights. Model weights are numerical parameters within an AI model that are adjusted during training to recognise patterns and features in the training data. A key finding was that the model itself did not store any of its training data, albeit a large part of its functionality was indirectly controlled by that training data.

The abandoned claims

During the trial, Getty abandoned two key claims. First, its claim relating to copyright infringement in relation to the training and development of the model, because there was no evidence that any of these acts had, in fact, taken place in the UK. Questions relating to training will however feature further in the ongoing US proceeding between the two companies (where, of course, the court will also consider arguments relating to fair use), and Getty has indicated in a press release that it would be taking "*forward findings of fact*" from the UK judgment in its US case. For now, this issue remains unconsidered by a UK court, and may therefore be significant in a different case.

Secondly, Getty abandoned its claim to copyright infringement arising out of alleged reproduction of its copyright works in outputs. Stability had, following the commencement of the proceedings, blocked the prompts alleged to have generated such infringing outputs, and Getty had in effect obtained its desired remedy. It remains possible that, in another case, a rights holder will have a much stronger incentive to pursue an "infringing output" copyright claim. Indeed, the court heard expert evidence on the concept of memorisation (in the context of the trade mark case), with both experts agreeing that images from the training data could be reproduced by a model in response to certain prompts as a result of overfitting, most likely because of data duplication during training. This may have more of a bearing in other situations (though here the evidence showed that in later versions of the model Stability had taken steps to remove images bearing watermarks from the training data, albeit the experts accepted that this would not 100% prevent such outputs appearing).

The decision

Trade mark infringement

Getty's trade mark infringement claim concerned the display of watermarks in outputs, which are identical or similar to its trade marks for GETTY IMAGES and iSTOCK. It argued that this amounted to trade mark infringement pursuant to sections 10(1), (2) and (3) of the Trade Marks Act 1994 (TMA).

To succeed in its trade mark infringement claim, Getty had to establish, as a threshold question and on the balance of probabilities, that Stable Diffusion had generated at least one output containing the relevant watermarks in the UK. The court made a number of model/version specific threshold findings, and concluded that there were some "*limited*" and "*historic*" instances of trade mark infringement pursuant to sections 10(1) and (2), but rejected the claim under section 10(3). The court decided not to give judgment on the passing off claims, given that this would not add further to the trade mark claims.

The trade mark claim involved the court in considering a number of interesting issues, including:

Getty's experimental prompts

Stability accepted that Stable Diffusion could be used to generate synthetic images including signs in the form of Getty's watermarks, but argued that Getty's evidence on this, in the form of experiments using different types of prompts, amounted to "*wilful contrivance*" and was not representative of real world users' prompts. Indeed, Stability argued that real world users would want to avoid generating images bearing watermarks.

The court found that there was no real world evidence that users would adopt some of the prompts used by Getty in its experiments (including prompts taken verbatim from captions on the Getty websites, re-worded

versions of those prompts, and prompts loosely inspired by other prompts). However, the court did find that certain invented prompts (including the words "*news photo*" and "*vector art*") were representative of prompts that real world users would use. On the facts, this meant that Getty only succeeded in relation to its iSTOCK marks and not in relation to its GETTY IMAGES marks.

Getty did also rely upon some real world examples generated "in the wild" using apparently random prompts. It could not establish these were generated in the UK, but the court accepted they were probative of the fact that such images would have been generated also in the UK.

The context of use

One particularly interesting issue was the extent to which the court should consider post-sale context of use, such as where more technically unsophisticated individuals might be exposed to the signs on synthetic images generated by someone else. This followed the guidance of the Supreme Court on post-sale confusion in trade mark infringement cases in *Iconix v DreamPairs* (discussed in [our article here](#)) which required an assessment of the realistic and representative way in which the average consumer will encounter the sign. The court suggested that it was not realistic and representative to conclude that there was a class of unsophisticated average consumer who would be shown "watermarked" images generated by a third party using Stable Diffusion. The position could be different if the trade mark was not a watermark, whose inclusion in outputs the court said would be considered undesirable by most users, but some other type of trade mark which would not be so undesirable (e.g. a popular character).

Use in the course of trade

Stability's use of the watermarks was found to be use of a sign in the course of trade because it amounted to using the sign for its own commercial communication, going beyond merely providing the technical conditions necessary for the service or storing the output images. Stability sought to pin control of the model regard on their users but the evidence demonstrated that, whilst users had some degree of control, they had no control over when a watermark was produced and in reality the appearance of the watermark was the result of the decisions Stability had made about its model rather than any prompting by the user. Accordingly, the only entity with any meaningful type of control was Stability. Again, this assessment of who has 'control' over the outputs of a model may have a bearing in future cases (including those relating to outputs that alleged to infringe copyright works). For example, the inclusion of other kinds of marks may in fact be actively sought in users' prompts.

Use in relation to goods and services

The court also found that Stability's relevant use of the sign was in relation to goods and services, which was liable to affect the functions of the trademark, in particular, the essential function of guaranteeing to consumers the origin of the goods. The threshold use of the watermarks was such as to create the impression that there was a material link in the course of trade between the goods concerned and Getty.

Identity and similarity of goods and services

The court found that providing access to Stable Diffusion was "*the provision of synthetic image outputs*" and found that that service was covered by certain goods and services in Getty's registrations in classes 9 and 41, but not in classes 16, 35 and 42. Getty's specifications in class 38 were found to cover Stable Diffusion in some but not all of its deployments.

Rejection of s.10(3) claim

Whilst the court found some limited infringement therefore under sections 10(1) and (2), it completely rejected the claim for infringement under section 10(3). The use it had found of the signs did not amount to detriment to distinctive character (dilution), detriment to reputation (tarnishment), or unfair advantage (free-riding). In

particular, in relation to the first two kinds of damage, Getty had not shown a change in the economic behaviour of the average consumer, or a serious likelihood of such a change.

Secondary copyright infringement

Getty's remaining copyright claim was one of secondary infringement said to arise from the importation (through downloading) and distribution (via Hugging Face) of Stable Diffusion in the UK. Its argument here was not that Stable Diffusion was *itself* a copy of Getty's copyright works, or that those works were stored within it. Rather, Getty alleged that Stable Diffusion was an "infringing copy" because the making of its model weights would have constituted copyright infringement if it had been carried out in the UK (pursuant to section 27(3) of the Copyright, Designs and Patents Act 1988 (CDPA)).

This part of the case raised two novel and important points of law:

- Was Stable Diffusion an article for the purposes of sections 22 and 23 of the CDPA?
- Was it an infringing copy for the purposes of section 27(3)?

Article?

There is no statutory definition of 'article' in the CDPA, and when it was introduced, storage media would, of course, merely have involved physical, tangible means of storage. However, the court agreed with Getty that it should adopt the 'always speaking principle' of statutory interpretation, and it therefore concluded that an article was not limited to tangible articles, and could be an electronic copy stored in an intangible medium (such as the AWS cloud).

Infringing copy?

However, the court went on to reject the broad interpretation Getty sought to attach to 'infringing copy'. Getty argued there was no requirement that an article must continue to retain a copy of the works and it was sufficient to show that the making of the model weights, had it been carried out in the UK, would have constituted an infringement of copyright. The court rejected this argument: an infringing copy had to be a 'copy', i.e. a reproduction of the copyright work. However, but the model weights of the various iterations did not store the visual information in the copyright works.

The court posed the central question: is an AI model which derives or results from a training process involving the exposure of model weights to infringing copies itself an infringing copy? It concluded that it was not. Whilst the model weights were altered during training by exposure to the copyright works, by the end of that process they did not store any of those copyright works. The model weights were "*purely the product of the patterns and features which they have learnt over time during the training process*".

The claim for secondary copyright infringement therefore failed.

Takeaways

The judgment has, perhaps unfairly, been described as a 'damp squib' by some commentators. Whilst it is narrower than it would otherwise have been given the claims that were abandoned, and rights holders will be disappointed at the conclusion on secondary copyright infringement, the court did make some important findings in relation to the trade mark claim. It is also possible that a different outcome may be reached on copyright infringement in a case with strong evidence of (real world) infringing outputs. For AI developers, meanwhile, the judgment underscores the value of robust filtering and guardrails to mitigate the risk of infringing outputs.

The decision also highlights how fact- and evidence-sensitive these types of claims are, particularly where, unlike trademark purchases in "traditional" IP infringement claims, each attempt to obtain evidence of infringement

will generate a different output, making reproducibility almost impossible. Furthermore, although Stability did not run a de minimis defence, such a defence may well feature in any future claims.

Subject to any appeal by either or both parties, all eyes will now turn to how the Government responds to its consultation in relation to copyright and GenAI (discussed in [our article here](#)). The question of where to strike the balance between the AI and creative industries (described by the judge as "*opposing factions*" when, in fact, Getty has launched its own AI tools) is a policy one. Despite the court's decision then, the creative industries will remain particularly focused in lobbying Parliament not to introduce provisions which will allow AI companies to scrape their copyright works to train their models, unless the relevant rights holder expressly opts their works out. This policy debate is set to take centre stage, with the [House of Lords hearing evidence in an inquiry](#) on AI and copyright on the very day the judgment was delivered.