



Neutral Citation Number: [2025] EWHC 2863 (Ch)

Case No: IL-2023-000007

IN THE HIGH COURT OF JUSTICE
BUSINESS AND PROPERTY COURTS OF ENGLAND AND WALES
INTELLECTUAL PROPERTY LIST (ChD)

Royal Courts of Justice
Rolls Building, Fetter Lane,
London, EC4A 1NL

Date: 04/11/2025

Before :

MRS JUSTICE JOANNA SMITH DBE

Between:

(1) **GETTY IMAGES (US) INC**
(a company incorporated under the laws of the State of
New York)
(2) **GETTY IMAGES INTERNATIONAL U.C.**
(a company incorporated under the laws of Ireland)
(3) **GETTY IMAGES (UK) LIMITED**
(4) **GETTY IMAGES DEVCO UK LIMITED**
(5) **ISTOCKPHOTO LP**
(a partnership incorporated under the laws of Canada)
(6) **THOMAS M. BARWICK, INC**
(a company incorporated under the laws of the State of
Washington)

Claimants

- and -

STABILITY AI LIMITED

Defendant

Lindsay Lane KC, Jessie Bowhill and Joshua Marshall (instructed by **Fieldfisher LLP**) for
the **Claimants**

Hugo Cuddigan KC, Edward Cronan and Henry Edwards (instructed by **Bird & Bird LLP**) for the **Defendant**

Hearing dates: 9-12, 17-20, 25-27 & 30 June 2025

**Approved Judgment
FOR PUBLICATION**

This judgment was handed down remotely at 10.00am on 4 November 2025 by circulation to the parties or their representatives by e-mail and by release to the National Archives.

.....

MRS JUSTICE JOANNA SMITH DBE

(A) INTRODUCTION	1
(B) FACTUAL BACKGROUND.....	3
Getty Images and their business.....	3
Stability and Stable Diffusion	6
(C) PROCEDURAL BACKGROUND.....	11
The Strike Out/Summary Judgment Application.....	11
Relevant Copyright Works.....	12
Identification of the Licensing Issue	12
Notices of Experiments	13
Procedural Orders made in respect of the Trade Mark Infringement Claim.....	14
Confidentiality.....	15
Applications made at trial	15
The Issues for Trial	15
Representation at trial	16
(D) THE WITNESSES AND EVIDENCE.....	16
Getty Images' Witnesses of Fact.....	16
Getty Images' CEA Notices	17
Stability's Witnesses of Fact	18
Stability's CEA Notices	20
The Technical Experts.....	21
Evidence of New York Law	22
The Trial Bundle	22
(E) LEGAL RESPONSIBILITY FOR STABLE DIFFUSION v1.X	22
Conclusion on legal responsibility for v1.x	31
(F) THE TRADE MARK INFRINGEMENT CLAIM.....	31
Introduction	31
The Threshold Issue: Incidence of synthetically generated watermarks*.....	32
Conclusion on the Threshold Issue	62
The Average Consumer.....	68
Context	75
SECTION 10(1) INFRINGEMENT	83
The Law	83
Application of the law to the facts	89
Conclusion on section 10(1) Infringement.....	119
SECTION 10(2) INFRINGEMENT	120
The Law	120
Likelihood of confusion: the global assessment	121
Conclusion on section 10(2) Infringement.....	124
SECTION 10(3) INFRINGEMENT	124
The Law	124
Getty Images' pleaded case.....	129
Change in Economic Behaviour.....	130
Detriment to Distinctive Character	132
Detriment to Reputation/Tarnishment.....	133
Unfair Advantage	142

Conclusion on section 10(3) Infringement.....	142
(G) PASSING OFF.....	142
(H) THE SECONDARY INFRINGEMENT CLAIM	143
The Statutory Framework	144
The key Issues between the parties	145
Preliminary points about the nature of Stable Diffusion.....	145
The Approach to Statutory Construction	149
The meaning of the terms “article” and “infringing copy”	152
The factual questions.....	162
(I) COPYRIGHT SUBSISTENCE AND OWNERSHIP	165
Legal Framework	165
The Disputed Works	167
(J) THE LICENSING ISSUE.....	171
Relevant Factual Background	172
Relevant Legal Principles: English law	172
Relevant Legal Principles: New York Law.....	176
Analysis of the outstanding issues of construction	179
(K) REMAINING OUTSTANDING ISSUES	196
Number of works used in training.....	196
Additional Damages	198
(L) CONCLUSION	199

Mrs Justice Joanna Smith DBE:

(A) INTRODUCTION

1. These are proceedings alleging primary and secondary copyright infringement, database right infringement, trade mark infringement and passing off against the Defendant (“**Stability**”), an **open-source** generative artificial intelligence (“**AI**”) company. The claim concerns Stability’s deep learning AI model (known as “**Stable Diffusion**” or “**the Model**”).
2. Shortly prior to closing submissions, the Claimants (collectively referred to as “**Getty Images**”) abandoned various aspects of their claim, thereby narrowing the issues to be determined by the court and rendering large parts of the opening submissions and evidence irrelevant. Nevertheless, the claim continues to raise issues of importance in the field of intellectual property in connection with the use of AI models such as Stable Diffusion.
3. **Deep learning** AI models use an ‘artificial neural network’ architecture designed to approximate the structure of synaptic connections in the brain. The neural network consists of multiple **layers** (hence the term **deep**) which create a hierarchical processing structure.
4. Stable Diffusion is a type of **generative AI** model known as a **diffusion model**, or more specifically a **latent diffusion model**. Broadly, it transforms an input (a user command or “**prompt**” in the form of written text or a “**seed**” image) into a desired output in the form of a synthesised image by modelling a probability distribution based on its training data and then **sampling** from that distribution. The development of a **stochastic** model of this type typically involves designing and building the architecture for the model which is then trained by repeated exposure to massive quantities of data, in this case in the form of human-generated digital images contained in datasets created by crawling and scraping images and associated descriptive captions from the Internet.
5. The model parameters (the “**model weights**” or “**biases**”) define the network connections in the model and are learnable parameters controlling the functionality of the network. Before training begins, the network’s weights are initialized with random values. As the network is exposed to the training data, the weights are **iteratively** updated to better satisfy an optimization criterion specified by engineers. Once the model is trained, running the network, referred to as **inference**, is (in simple terms) an input-output system in which the user specifies inputs, the trained network performs computations on those inputs and then generates the desired output.
6. Although there are differences between the various versions of Stable Diffusion, essential to these models is a process which starts with a random noise image. The trained network (conditioned by the user-specified prompt) iteratively removes the noise so as to create an image which is semantically consistent with the prompt.
7. Inference does not require the use of any training data and the model itself does not store training data. However, a large part of its functionality is indirectly controlled via the training data. In other words, the way in which the network makes use of its multiple layers is the result of the training process.

8. The claim as pleaded seeks to protect what Getty Images describe as “the lifeblood” of their business, namely millions of high-quality photographic images of, amongst other things, world events, sporting moments, celebrities, architecture, nature and travel (“**the Visual Assets**”), created over many years by hundreds of thousands of photographers. At the heart of that case is the allegation that, in order to develop and train Stable Diffusion, Stability has scraped millions of Visual Assets (a substantial proportion of which are said to comprise original artistic works and/or film works owned by, or exclusively licensed to, the First Claimant in which copyright subsists (“**the Copyright Works**”) from Getty Images’ websites without consent and used those images unlawfully to train and develop a number of versions of Stable Diffusion.
9. Notwithstanding their pleaded case, it is now acknowledged by Getty Images that (i) there is no evidence that the training and development of Stable Diffusion took place in the United Kingdom (such that what has been called “**the Training and Development Claim**” has been abandoned); (ii) the prompts which it was alleged had been used to generate the examples of infringing output from the Model in evidence in these proceedings have been blocked by Stability such that the relief to which Getty Images would have been entitled in respect of their allegations of primary infringement of copyright (referred to as “**the Outputs Claim**”) has now been substantially achieved. Thus the Outputs Claim has also been abandoned; and (iii) given its inherent link to the Training and Development Claim and the Outputs Claim, a claim for database rights infringement (“**the Database Rights Infringement Claim**”) can now no longer be advanced.
10. However, Getty Images continue to advance a case that normal use of Stable Diffusion by users in the United Kingdom will, in some cases, generate synthetic images bearing Getty Images’ own trade marks, contrary to sections 10(1), 10(2) and 10(3) of the Trade Marks Act 1994 (the “**TMA**”) (“**the Trade Mark Infringement Claim**”), and/or constituting an actionable misrepresentation under the law of passing off (the “**Passing Off Claim**”). They also maintain a case that, contrary to sections 22 and 23 of the Copyright, Designs and Patents Act 1988 (the “**CDPA**”), Stability has imported into the UK, otherwise than for private and domestic use, possessed in the course of business, sold or let for hire or offered or exposed for sale or hire, or distributed an article, namely Stable Diffusion, which is and which Stability knew or had reason to believe is an infringing copy of the Copyright Works.
11. Getty Images do not say that Stable Diffusion is itself a copy of, or that it stores within it any copies of, the Copyright Works. However, pursuant to section 27(3) CDPA, Getty Images contend that Stable Diffusion is an infringing copy under the CDPA because the making of its model weights would have constituted infringement of the Copyright Works had it been carried out in the UK (“**the Secondary Infringement Claim**”).
12. Both sides emphasise the significance of this case to the different industries they represent: the creative industry on one side and the AI industry and innovators on the other. Where the balance should be struck between the interests of these opposing factions is of very real societal importance. Getty Images deny that their claim represents a threat to the AI industry or an attempt to curtail the development and use of AI models such as Stable Diffusion. However, their case remains that if creative industries are exploited by innovators such as Stability without regard to the efforts and

intellectual property rights of creators, then such exploitation will pose an existential threat to those creative industries for generations to come.

13. As to whether this judgment will, in reality, have anything to say on the balance to be struck between the two warring factions, it is worth observing at the outset that this court can only determine the issues that arise on the (diminished) case that remains before it having regard to the available evidence and the arguments advanced by the parties. It is no part of this court's task to consider issues that have been abandoned or to consider arguments that are no longer of relevance to the outstanding issues.
14. Attached to this Judgment at Appendix A is a Glossary of the key technical terms used in this Judgment when discussing the technology. Those terms have also been emboldened in the text of the Judgment for ease of reference.

(B) FACTUAL BACKGROUND

Getty Images and their business

15. I do not understand the following description of Getty Images and their business, much of which is set forth in Getty Images' Opening Submissions, to be controversial.
16. The Getty Images business was founded in 1995 through the incorporation of Getty Communications plc and, since then, has grown through a series of mergers and acquisitions, including the acquisition of the Canadian company iStockphoto, Inc. The First to Fifth Claimants are members of the Getty Images group ("**the Group**"). The First Claimant was incorporated under the laws of New York, the Second Claimant was incorporated under the laws of Ireland, the Third and Fourth Claimants were incorporated under the laws of England and Wales and the Fifth Claimant was incorporated under the laws of Canada.
17. The Group is now a pre-eminent global visual content creator and market-place. Its business involves the licensing of millions of Visual Assets including photographs, video footage and illustrations, as well as audio assets (collectively "**Content**"), to individuals and business users, such as newspapers, magazines, production companies, advertising agencies, banks, airlines, insurance companies and pharmaceutical companies, in more than 200 countries worldwide. Getty Images licenses the Content in a variety of ways to end users via their standard licensing agreements and, in some cases, bespoke customer licensing deals.
18. The Copyright Works are said to make up a substantial proportion of the Content. Getty Images assert copying during the training and development of the Model by Stability in respect of millions of Copyright Works of which the First Claimant is either the owner or exclusive licensee. As I have said, it is now accepted that the training and development of the Model did not take place in the United Kingdom and so this is not an allegation that I need to address.
19. The Content exists in a sophisticated curated database ("**the Getty Images Database**") with a plethora of associated metadata which includes, amongst other things, the content type, the date the content was captured or created, the pixel and file size, information relating to the creator of the Content, relevant keywords and, importantly, the relevant caption. Captions can vary from being a lengthy detailed description of the

Content to a short generic description. There are two main categories of Content: (i) editorial content comprising Content that is newsworthy or of public interest and depicts real life people, places and events; and (ii) creative content, comprising pre-shot stock photography, illustrations and video used for a variety of purposes.

20. The Content is said by Getty Images to be highly desirable for use in connection with AI and machine learning because of its high quality and because it is accompanied by content specific, detailed captions and rich metadata.
21. Getty Images make the Content available through websites at gettyimages.com (launched in 2001), gettyimages.co.uk, and istockphoto.com (acquired in 2006 (the “**iStock Website**”) (collectively the “**Getty Images Websites**”). The Getty Images Websites comprise hundreds of millions of Visual Assets, together with associated captions covering a broad range of subject matter. Customers are able to browse the Getty Images Websites using keywords and filters to locate their desired Content, including the associated metadata and captions. The iStock Website is searchable and contains a library of pre-shot creative content with accompanying captions, including “**vector files**”, namely a resolution independent depiction of a visual image that can be scaled up or down and is often used for illustrations. Each Visual Asset that is available through the Getty Images Websites has an associated page that contains a unique uniform resource locator (“**URL**”) pointing to a location where the image is stored, together with an “alt text” tag containing a caption for the image. The Getty Images Websites are hosted by servers and are available via both iOS and Android apps, in 23 different languages.
22. The Content is acquired by Getty Images in a variety of different ways: (i) by acquiring the Content outright, for example by way of an assignment from a rightsholder or the acquisition of a company with a substantial portfolio of Content; (ii) by entering into licence agreements with third party photographers and copyright owners; (iii) via photographers and videographers who are employed by a member of the Group in various jurisdictions; and (iv) through contracts with ‘stringer’ photographers who are hired to cover a specific event and paid a day rate, with the terms of the contract assigning copyright to Getty Images or granting it a perpetual exclusive licence to the Content. Many of these contracts and agreements take the form of a Unified Contributor Agreement (“**Contributor Agreement**”) or an iStock Artists Supply Agreement (“**iStock ASA**”). These are template agreements which have been updated and varied over time and which govern the terms on which the Content is licensed to Getty Images.
23. The Sixth Claimant is an example of a creative Content contributor. It is a US company founded in June 2005 and has an exclusive arrangement with Getty Images whereby it produces commercial imagery for Getty Images to license via the Getty Images Websites. The Sixth Claimant is the trading vehicle of the photographer and full-time employee and director, Thomas Barwick (“**Mr Barwick**”). It has approximately 35,000 still assets and 15,000 video assets available on the Getty Images Websites. The Sixth Claimant has entered into various iterations of both the Contributor Agreement and the iStock ASA with the First Claimant, the most recent being a Contributor Agreement dated 3 October 2023.
24. Getty Images own various trade marks (together the “**Marks**”) which relate to the Getty Images and iStock name and logo:

i) The First Claimant is the registered proprietor of three UK registered trade marks (“**the Getty Images Marks**”):

- a) word mark UKTM No. UK00911410859 (“**UK859**”) registered on 10 December 2012¹ for GETTY IMAGES in respect of goods and services falling within classes 9, 42 and 45;
- b) word mark UKTM No. UK00902313005 (“**UK005**”), registered on 24 July 2001 for GETTY IMAGES in respect of goods and services falling within classes 9, 16, 35, 38, 39, 40, 41 and 42; and
- c) figurative word mark UKTM No. UK00908257925 (“**UK925**”), registered on 29 April 2009 for:
in respect of goods and services falling within classes 9, 16, 35, 38, 39,
gettyimages
41, 42 and 45.

ii) The Fifth Claimant is the registered proprietor of two UK registered trade marks (the “**ISTOCK Marks**”):

- a) word mark UKTM No. UK00908257297 (“**UK297**”), registered on 29 April 2009 for ISTOCK in respect of goods and services falling within classes 16, 40 and 41; and
- b) word mark UKTM No. UK00906776819 (“**UK819**”) registered on 25 March 2008 for ISTOCK in respect of goods and services falling within classes 9, 35, 38, 42 and 45.

25. It is common ground that the Marks are inherently distinctive and have acquired a reputation by virtue of the extensive use made of them by Getty Images. The Getty Images logo and brand have remained largely unchanged and in constant use for the last 30 years, since the launch of Getty Images in 1995. The ISTOCK Marks have existed since at least 2003. The Marks are used in all aspects of Getty Images’ business such as trading names, websites, social media, company letterheads, building signage, email signatures and all marketing materials and merchandise. Unchallenged evidence from Getty Images shows a high volume of followers, interactions, impressions and reach on the Getty Images social media accounts during the period 2020-2022. Furthermore, editorial content is used in much of the world’s media, and includes a photo credit on any image that is used, which refers to Getty Images (and sometimes the name of the photographer). Getty Images undertakes extensive marketing, spending very significant sums every year. Consistent with the breadth and depth of its operations, in each of the years 2017-2022, the First, Second, Third and Fifth Claimants’ UK revenue generated under the Getty and ISTOCK brands has run into many millions of pounds.

26. Each of the Visual Assets that appears on the Getty Images Websites displays a watermark that contains one or other of the Marks: either a Getty Images watermark

¹ All registration dates are taken from the date of filing of the application for registration in accordance with section 40(3) Trade Marks Act 1994.

containing a Getty Mark or, on the iStock Website, an iStock watermark containing an ISTOCK Mark. It is Getty Images' case that these watermarks have become iconic in their own right. If a Visual Asset is downloaded from the Getty Images Websites, it will feature the Getty Images or iStock watermark with the Marks appearing within a grey translucent banner which is overlaid on the image. In the case of still photographs, the name of the photographer will appear beneath the Mark in the watermark, as can be seen in the following example, taken by Mr Barwick:



27. It is only by properly licensing the Visual Asset from Getty Images that a version is made available to the customer without the Getty Images or iStock watermark on it.
28. In 2023, Getty Images launched its own AI software tools, developed in conjunction with NVIDIA, "Generative AI by Getty Images" and "Generative AI by iStock", which are available to Getty Images' customers via subscription through an application program interface ("API") on the Getty Images Websites ("the GAI"). The GAI was trained on Getty Images' creative pre-shot content library (which is all licensed content).

Stability and Stable Diffusion

29. Stability was incorporated on 4 November 2019 in England and Wales. It carries on business in the field of machine learning software, including **Deep Learning** models for image and music generation, and large language models ("LLMs") for the generation of text output.
30. Stable Diffusion is based on independent research work undertaken by academic researchers (including Professor Björn Ommer ("Professor Ommer") and Mr Robin

Rombach (“**Mr Rombach**”) at the Computer Vision and Learning Group (“**CompVis**”) at Ludwig Maximilian University of Munich (“**LMU**”) and IWR Heidelberg University, Germany). Mr Rombach co-authored an academic paper entitled “High Resolution Image Synthesis with Latent Diffusion Models” (“**the Latent Diffusion Paper**”), first published on 20 December 2021 and subsequently published in revised form on 13 April 2022. Mr Patrick Esser (“**Mr Esser**”) of Runway ML (“**Runway**”) was also involved in the publication of the Latent Diffusion Paper. This paper proposed a new method of generating images using a diffusion model which was trained on images transformed using existing trained **autoencoders** to a reduced-definition latent representational space. Such a latent diffusion model offered technical advantages over existing (i.e. pixel-based) diffusion models including being less resource intensive and able to generate higher resolution output images due to its efficiency.

31. On 21 December 2021, CompVis published source code and pre-trained model weights for its own Latent Diffusion model via a CompVis public web portal called **Github**, which allows users to upload and share code and data. In April 2022, CompVis published an updated version of this model on the CompVis Github page, trained by members of CompVis, including Mr Rombach. It is common ground that Stability had no responsibility for this model. It is, however, the unchallenged evidence of Stability’s technical expert in these proceedings that the Latent Diffusion Paper, together with the materials made available on the CompVis Github page, provided the structure and code for the “underlying model architecture” of the first iteration of Stable Diffusion.
32. Stable Diffusion was originally released further to an agreement between Mr Rombach and Mr Emad Mostaque (“**Mr Mostaque**”), founder and CEO of Stability, pursuant to which Stability gave Mr Rombach and another CompVis researcher access (via the internet) to cloud hosting and processing resources (“**the AWS Cluster**”) made available to Stability by Amazon Web Services Inc (“**AWS**”). The AWS Cluster was located outside the United Kingdom. Stability says that it utilised the AWS Cluster with the aim of offering such services to academic and other non-profit researchers and so to promote the development and growth of **open source** machine learning models. An article on the LMU website dated 1 September 2022 explains that “[i]n their project, the LMU scientists had the support of the start up Stability.Ai, on whose servers the AI model was trained”. Professor Ommer describes how “[t]his additional computing power and the extra training examples turned our AI model into one of the most powerful image synthesis algorithms”.
33. On 10 August 2022, Stability announced the first stage of the release of Stable Diffusion to researchers on its website, describing it as “a text-to-image model empowering billions of people to create stunning art within seconds”. The announcement went on to say that Stable Diffusion “runs on under 10GB of **VRAM** on consumer **GPUs**, generating images at 512x512 pixels in a few seconds”². It was intended that Stable Diffusion would shortly be made available to the public (i.e. on an open source basis), thereby “democratizing image generation”.
34. Thereafter, Stable Diffusion was released to the public in various iterations or ‘**checkpoints**’, namely v1.1, v1.2, v1.3 and v1.4. These versions (published together as “**v1.x**”) were made available for download (including in the United Kingdom) by

² A VRAM or Video Random Access Memory is a dedicated memory on a computer’s GPU or Graphics Processing Unit that stores and manages data related to graphics and video processing.

release of the source code and model weights via the CompVis GitHub and CompVis Hugging Face web portals on or around 22 August 2022. A user accessing the Model via these portals is able to download the inference code from GitHub and the model weights from Hugging Face and set up his or her local computer to run the software, thereby running the inference offline (“**the Direct Download Method**”). Alternatively, a user can access the Hugging Face ‘Diffusers’ library which contains inference code and allows the user to download the model weights through a code interface (“**the Diffusers Method**”). This requires use of additional software including MiniConda, the Nvidia Cuda Toolkit and PyTorch. Once the Model has been downloaded using either of these methods, expert users have the opportunity to modify the code and to run additional training with their own data.

35. At the same time, Stability announced the “Stable Diffusion Public Release” on its website, providing a link to the **Model Card** and weights and also recommending the use of “**DreamStudio**”, a commercial platform hosted outside the UK which enables users in the UK and elsewhere to access Stable Diffusion v1.4 without the need to download the Model; users are able to run the inference on Stability’s computing system using a web interface. Access to DreamStudio was provided by Stability via a public test version of DreamStudio referred to as **DreamStudio Beta**, available via Stability’s website and also accessible at beta.dreamstudio.ai. The Model Cards for v1.1, v1.2, v1.3 and v1.4 explain, amongst other things, that (each version of) the Model is a Latent Diffusion Model that uses a fixed pretrained text encoder (CLIP ViT-L/14) and (under the heading “Limitations”) that it has been trained on a “large-scale dataset LAION-5B which contains adult material”.
36. Stability’s announcement on 22 August 2022 anticipated the continuing release of optimized versions of the Model together with the provision of **API** access³. A further, unofficial, version of the Model (v1.5) was released in or around October 2022. Although originally forming part of Getty Images’ pleaded case, references to this version have been excised from the pleadings and (although it is mentioned in some of the evidence) it no longer forms part of that case.
37. At some time in the Autumn of 2022, after the release of Stable Diffusion v1.x, Mr Rombach joined Stability as an employee and Head of Research. He was based in Germany.
38. A TechCrunch article published on 17 October 2022 quotes Mr Mostaque as saying that DreamStudio had more than 1.5 million users “who’ve created over 170 million images”. The article also states that, according to a Stability press release, the open source version of Stable Diffusion had been downloaded more than 200,000 times.
39. On 24 November 2022, Stability launched Stable Diffusion 2.0 (“**v2.0**”), making available the Python source code for v2.0 on Stability Github and the pre-trained model weights on Stability Hugging Face on an open source basis. Stability AI is named on the licence (CreativeML Open RAIL++-M) as a copyright owner. The Stability Github page explains that v2.0 has been “trained from scratch”. In other words, it did not use any of the model weights obtained from the process of training the version 1.x Models, but represented an entirely fresh start, using a new training dataset and a different **text**

³ An API, or Application Programming Interface is a set of rules and specifications that allow different software systems to communicate with each other.

encoder model. Stability released an API Platform for v2.0 on 25 November 2022, thereby enabling users to run the inference on Stability's computing network using an API without the need to download the Model. Stable Diffusion v2.0 was also made available on DreamStudio from around 1 December 2022 until early 2024.

40. Version 2.1 (“**v2.1**”) of Stable Diffusion was released in December 2022 as a further development of the existing v2.0 checkpoint, using additional steps but tweaking the settings of the NSFW Filter. Further checkpoints of v2.0 are referred to together as **v2.x**. The Model Cards for v2.x Models explain amongst other things, that (each version of) the Model is a Latent Diffusion Model that uses a fixed pretrained text encoder (OpenCLIP-ViT/H) and (under the heading “Limitations”) that the Model was trained “on a subset of the large-scale dataset LAION-5B, which contains adult, violent and sexual content”. In contrast to the v1.x Model Cards, the v2.0 Model Card goes on to say that “[t]o partially mitigate this, we have filtered the dataset using LAION’s NSFW detector…”, which is later said to produce a “p_unsafe score of 0.1 (conservative)”. The acronym NSFW is short for **Not Safe For Work** (“**NSFW**”) material.
41. On 13 April 2023, Stability announced the release of Stable Diffusion XL Beta and thereafter made available further model checkpoints in this series (together “**SD XL**”), capable of outputting higher resolution images of 1024x1024 pixels. Thus in June, July and November 2023, Stability released Stable Diffusion **SD XL 0.9**, **XL Base 1.0** and **XL Turbo**. SD XL Beta was made available through an API and the DreamStudio platform, while the model weights and associated source code for versions SD XL 0.9, XL Base 1.0 and XL Turbo were made available for download on Hugging Face (and in one case on Git Hub). SD XL Turbo makes use of a technique called adversarial diffusion distillation (“**ADD**”) and combines aspects of both generative adversarial networks (“**GANS**”) and diffusion models to create high quality images quickly.
42. The SD XL Model Cards explain amongst other things, that (each version of) the Model is a Latent Diffusion Model that uses two fixed pretrained text encoders (OpenCLIP-ViT/G and CLIP-ViT/L). The SD XL Model Card explains that it is a distilled version of SD XL 1.0 and that it is based on the novel training method called ADD. These Model Cards no longer record how these Models were trained and, although they continue to identify “Limitations”, they do not say anything about the presence of adult, violent or sexual content in the training dataset.
43. Each of the Model Cards for v1.x, v2.x and SD XL identifies “Out-of-Scope Use” in the following terms: “[t]he model was not trained to be factual or true representations of people or events and therefore using the model to generate such content is out-of-scope for the abilities of this model”.
44. On 18 July 2023, Stability launched the Developer Platform API (“**the Developer Platform**”) at <https://platform.stability.ai> designed as a reboot of the prior API Platform. This platform makes APIs available to subscribers, thereby enabling subscribers’ applications to request services from various versions of the Model, remotely hosted outside the UK by Stability. It does not enable subscribers to download model weights or associated source code but instead enables them to run inference on AWS. Originally Stable Diffusion v1.x and v2.x variants were made available through APIs but support for these variants has since been discontinued.

45. In addition to the Stable Diffusion XL Models, Stability also developed a further Stable Diffusion Model known as Stable Diffusion 1.6 (“**v1.6**”). This Model is based on the architecture of Stable Diffusion XL and is optimized to generate images which are 512x512 in resolution. V1.6 was made available for use in around November 2023. The model weights for v1.6 have not been published and it would appear that v1.6 is not available to download. The Developer Platform now facilitates access to SD XL and to v1.6.
46. It is common ground that all versions of Stable Diffusion were trained using various subsets of the LAION-5B dataset (“**LAION-5B**” and “**LAION-Subsets**”) assembled by LAION e.V. (“**LAION**”), a not-for-profit organisation registered in Hamburg, Germany. LAION researchers first announced the publication of a dataset known as LAION-400M (“**LAION-400M**”) on 20 August 2021. LAION-400M, as its name suggests, comprised approximately 400 million **CLIP-filtered** (Contrastive Language-Image Pre-training) **URL-text** pairs⁴ (URLs with associated alt-text captions) and was created by LAION by filtering the Common Crawl public web archive for images and storing these alongside their HTML alt-text (hidden text associated with the image on a web page). The Latent Diffusion model developed by CompVis was trained on LAION-400M.
47. LAION researchers announced the publication of LAION-5B and the LAION-5B Subsets on 31 March 2022. Stability admits that at or around this time, it donated support to LAION in the form of hosting services comprising access to the AWS Cluster. The available evidence suggests that LAION-5B was created in a similar manner to LAION-4400M, but with the addition of further filtering steps. LAION-5B comprises metadata including 5.85 billion URL-text pairs. The LAION-Subsets (said by Getty Images to contain millions of Copyright Works) were produced by filtering LAION-5B against specific sets of requirements and together comprise approximately 3 billion URL-text pairs from the LAION 5B dataset. Information contained in the Model Cards for the various iterations of Stable Diffusion indicates that these Models were trained on LAION-Subsets which included LAION-2B-en, LAION-improved-aesthetics, LAION-aesthetics v2 5+, LAION-high-resolution and LAION-A. The Model Card for v2.0 gives a ‘shout out’ to “The DeepFloyd team at Stability AI, for creating the subset of LAION-5B dataset used to train the model”.
48. It is common ground that for training purposes it would have been necessary to download the images from the URLs in the LAION-Subsets – a process known as **materialisation**. Stability says that the training process involved downloading and storing copies of each image obtained from the URLs in the relevant dataset on Amazon Simple Storage Service (“**Amazon S3**”) on the AWS Cluster, retrieving those images and then making temporary copies of them in the VRAM of the GPUs performing the training on the AWS Cluster. Owing to the abandonment of the Training and Development Claim, Getty Images no longer seek to advance a case that this training process took place in the UK. Stability’s case has always been that training took place on the AWS Cluster outside the UK, but it has sought no declaration to this effect and accordingly I need make no finding on the point.

⁴ URLs or Uniform Resource Locators are hyperlinks (via a web address) to the original internet location at which the referenced image is made available to the public by the website operator.

49. Stability accepts that “at least some” LAION-Subsets contain URLs referencing images on the Getty Images Websites and that “at least some” images from the Getty Images Websites were used during the training of Stable Diffusion. However, there is a dispute over whether this court should make any findings as to the number of images from the Getty Images Websites which were so used, it being Stability’s case that the particular images used will depend on the starting dataset and the filters applied to it for each training run. Filters, such as the NSFW filter to which I have already referred, may be applied to the training dataset to remove (or attempt to remove) undesirable images. I shall return to this in due course.
50. Stability also accepts that Stable Diffusion *may* be used to generate synthetic images which include the Marks in the form of Getty Images’ watermarks. However, it contends, broadly, that (i) where such images are generated by a user, this is the result of third party use of Stable Diffusion and not a statement or commercial communication attributable to Stability, or for which Stability is responsible in law; (ii) any such generation of watermarks does not amount to use of any sign in those watermarks in the course of trade; and that (iii) “watermarked” synthetic image outputs will only be generated with wilful contrivance of the user.

(C) PROCEDURAL BACKGROUND

51. The claim in these proceedings was issued on 16 January 2023 (prior to the launch of SD XL) and the pleadings on both sides have since been the subject of numerous amendments⁵, clarification by way of Further Information, and addition in the form of Statements of Case on various discrete issues.
52. On 22 April 2024, Master McQuail ordered a first trial to determine liability.
53. Owing to the many complex procedural issues arising (themselves often a function of the novelty of the factual and legal issues in this case), it proved necessary for the court to adopt almost unprecedented levels of case management involving frequent substantial hearings taking place in the 6 months preceding trial. Many of these hearings resulted in a reported judgment and it is therefore unnecessary for me to say much about the procedural history beyond making the following points which are relevant to an understanding of this judgment and the issues that arise for determination.

The Strike Out/Summary Judgment Application

54. On 28 July 2023, before the filing of its Defence, Stability issued an application for strike out and/or reverse summary judgment (amongst other things) in respect of the Training and Development Claim and the Secondary Infringement Claim. I dismissed the application for summary judgment ([2023] EWHC 3090 (Ch)) but I observe that, as will already be apparent, the Training and Development Claim has not survived the evidence at trial.
55. Although I considered one of the (two) main legal arguments now arising on the Secondary Infringement Claim in my judgment (namely the true statutory interpretation of the word “article” in section 27 CDPA), I ultimately declined to decide the point.

⁵ References to statements of case in this judgment will be to the most recent version unless the contrary is stated.

Instead I made it clear that the point of law raised by Stability must be decided at trial further to full and comprehensive argument from both sides and further to factual findings as to the nature of the acts said to give rise to the secondary infringement. In so far as I expressed views on the authorities relevant to this point of law in the judgment, I do not consider myself bound by what were essentially preliminary views expressed in the absence of full argument on the point.

Relevant Copyright Works

56. Getty Images' Particulars of Claim asserts the copying by Stable Diffusion of millions of Copyright Works of which the First Claimant is either the copyright owner or the exclusive licensee pursuant to exclusive licence agreements set out in Annex 3A thereto. In the interests of proportionality, Getty Images chose to rely at trial upon eleven "**Sample Works**" (identified as **Works A-K**) for the purposes of establishing subsistence and ownership of copyright. Sample Works A-D are each artistic works, the copyright in which is alleged to be owned outright by the First Claimant; Sample Works E-H are artistic works (namely photographs) created by Mr Barwick using his own intellectual creativity and Sample Works I-K are films of which Mr Barwick was the principal director and producer. Getty Images alleges that the artistic copyright in Works E-K is owned by the Sixth Claimant and exclusively licensed to the First Claimant.
57. On 6 September 2024, Getty Images filed and served their Statement of Case on Infringement ("SOCI"). In the SOCI, Getty Images identified seventeen Copyright Works A1-A17 ("**the SOCI Works**") and then went on to provide particulars of subsistence and ownership in respect of those Copyright Works, providing the identity and employment status of the author of each work, the date of the work and (where the author of the work was an exclusive licensor of the work) the exclusive licence agreement relied upon by Getty Images to establish legal entitlement on the part of the First Claimant.
58. For the purposes of the copyright infringement claims originally pleaded, including the Secondary Infringement Claim, the parties agreed that certain specific issues should be determined by reference to identified Copyright Works. In the recitals to an Order of 17 January 2025 ("**the January 2025 Order**"), it was recorded that Stability had confirmed for the purposes of these proceedings that it would not seek to challenge the case advanced by Getty Images that Sample Works A1-A17 were used in the training of v1.4, v2.0 (and all its sub-versions) and SD XL (and all its sub-versions). In respect of the Secondary Infringement Claim, the January 2025 Order provided that the issues of subsistence and ownership of copyright "shall be decided on the basis of SOCI Works A1-A17 and Sample Works A-K". I shall return to consider the outstanding issues of entitlement arising in respect of these works (the so-called "**Ownership Issues**") when I address the Secondary Infringement Claim.

Identification of the Licensing Issue

59. As I recorded in a judgment handed down on 14 January 2025 ([2025] EWHC 38 (Ch)), it is Getty Images' case that there are likely to be in excess of 50,000 photographers and Content contributors who are owners of the relevant copyright subsisting in Content that has been licensed on an exclusive basis to the First Claimant over several decades. By reason of these exclusive licences, each of these copyright owners is alleged to have

a concurrent right of action with the First Claimant under sections 101 and 102 CDPA to restrain Stability's unlawful acts of copyright infringement and to seek relief in respect of such acts.

60. Against that background, the allegations of copyright infringement were originally brought by the Sixth Claimant in a representative capacity for all copyright owners who had granted Getty Images an exclusive licence in relation to the Copyright Works in issue in the proceedings. However, following a successful challenge by Stability to this representative claim ([2025] EWHC 38 (Ch)), Getty Images was instead permitted to proceed under CPR 19.3(1) and s.102(1) CDPA without joining any of the "Exclusive Licensors" (a term defined in the recitals to the January 2025 Order as meaning, in summary, any copyright owner jointly entitled to the remedies claimed by the First or Second Claimant in these proceedings as a result of being a party to any of the exclusive licence agreements set out in Schedule 1 to that Order). The January 2025 Order provided that this liability trial would "determine whether the Schedule 1 Agreements are exclusive licences within the meaning of s.92 CDPA (**"the Licensing Issue"**) by reference to a sample...".
61. By a further Order of 27 February 2025 ("**the February 2025 Order**"), the Court gave various sampling directions for trial, including that:
 - i) in relation to the Licensing Issue, the liability trial would proceed on the basis of 14 sample licence agreements, namely licence agreements #2, #3, #8, #10, #11, #13, #17, #19, #30, #32, #33, #36, #37 and #38 ("**the Sample Licences**"). These had been selected by the parties from 14 groups of licence agreements identified in Schedule 1 to the February 2025 Order as Groups A to N;
 - ii) the Court's findings in relation to each of the relevant Sample Licences would be applicable to all licence agreements falling within the respective groups A to N to which the relevant Sample Licence relates; and
 - iii) all licences asserted as relevant by Getty Images to the Sixth Claimant and/or the licensors of SOCI Works (set out in Schedule 2 to the February 2025 Order) would be considered at trial in any event.
62. The numbering of the Sample Licences set out in Schedule 1 to the February 2025 Order corresponds to numbered licence agreements set out in Annex 3A to the Particulars of Claim, which lists all of the licence agreements which are in issue in these proceedings and on which Getty Images rely. Those agreements also capture the relevant licence agreements for each of the exclusively licensed SOCI Works (i.e. A14-A16). For the purposes of this claim, Getty Images only relies upon licence agreements which contain a "right to control claims" clause.
63. The Licensing Issue, which requires consideration of New York Law in connection with the true interpretation of the Sample Licences, falls to be considered in the context of the Secondary Infringement Claim only.

Notices of Experiments

64. On 23 December 2024, the parties exchanged their respective Notices of Experiments ("NoE") and on 24 January 2025 they each served a Reply. On 5 March 2025, Stability

served a Notice of Reply Experiments, to which Getty Images replied on 18 March 2025.

65. By the February 2025 Order, the Court permitted the parties to rely upon their experiments as set out in their respective NoEs (including Stability's Notice of Reply Experiments) for the purposes of establishing by experimental proof the facts set out therein. In this judgment I shall refer to the section of Getty Images' NoE that addresses the generation of images which may contain watermarks as "**the Getty Watermark Experiments**". I shall refer to Stability's NoE and its Notice of Reply Experiments as "**the Stability Watermark Experiments**".

Procedural Orders made in respect of the Trade Mark Infringement Claim

66. The eighth recital to an Order of 1 November 2023 records Getty Images' confirmation that they do not rely upon synthetic outputs which are generated in response to users entering "prompts that contain signs corresponding to the Trade Marks of the Claimants" as acts of trade mark infringement or passing off.

67. The Trade Mark Infringement Claim is supported by two Annexes attached to the Particulars of Claim, namely **Annex 8H** and **Annex 8I**. Annex 8H contains examples of synthetic images bearing the Getty Images and/or ISTOCK watermark which were generated by the Getty Watermark Experiments. Annex 8I contains synthetic images bearing the Getty Images and/or ISTOCK watermark which were generated "in the wild" by third parties.

68. At the PTR on 22 May 2025 Getty Images were ordered to select four images, two from Annex 8H and two from Annex 8I (at least one for Getty Images and at least one for ISTOCK). These examples were to be used "for the purposes of the parties' submissions as to the comparison between mark and sign for trade mark infringement". This was not otherwise to restrict the scope of the parties' submissions on the Trade Mark Infringement Claim. Stability declined the opportunity to be involved in the selection process of the images to be used for the purposes of the comparison.

69. By reason of an Order dated 9 July 2024 (amended by a Consent Order dated 7 February 2025), the issue of whether Stable Diffusion generates synthetic images that bear the Marks in the form of watermarks is to proceed on the basis of the following models and access mechanisms:

- i) Stable Diffusion v.1.1-1.4, v2.1 and SD XL 1.0 as made available for download via Hugging Face; and
- ii) Stable Diffusion v1.6 as made available for use via DreamStudio.

70. This was designed to obviate the need for Getty Images to conduct experiments and/or produce images for every sub-version of the Model using every access mechanism. The Getty Watermark Experiments used these access mechanisms for experiments on versions 1.2, 1.3, 1.4 and 2.1. For experiments on v1.6, Getty Images accessed the code in the United States via an API available at <https://platform.stability.ai/>.

Confidentiality

71. On 10 October 2023 Master McQuail made a Confidentiality Order establishing a confidentiality regime at the request of the parties for the purposes of protecting confidential information disclosed in these proceedings. The existence of this regime meant that (in some cases) evidence on which the parties relied at trial was designated Confidential or Highly Confidential and that it was necessary to sit in private for short periods during the cross examination of witnesses in order to protect this confidentiality. I made it clear to the parties, however, that as much of the trial as possible must be in public in order to secure the proper administration of justice.
72. During the course of the trial, it became clear that the net of confidentiality had in some instances been cast too wide and I required both parties to review their evidence and disclosure with a view to de-designating material that should not have been caught in the original net. Shortly after close of trial the parties provided me with a further bundle of witness statements and exhibits with updated confidentiality markings to reflect de-designations.
73. Upon this judgment being provided to the parties in draft, they considered whether continuing confidentiality in material that is referred to in the judgment is necessary, having regard to the provisions of CPR 39.2. With two very minor redactions requested by Stability (which I am satisfied are necessary in respect of information which is in the nature of a technical trade secret) this judgment is being published in full, consistent with the principles of open justice.

Applications made at Trial

74. On the first day of trial, Stability sought clarification as to whether an allegation in the Particulars of Claim to the effect that Stable Diffusion “can be used to create images that contain pornography, violent imagery and propaganda” and that “[a]ny association with such content will tarnish the reputation of the Trade Marks” could properly be read as including reference to Child Sexual Abuse Materials (“CSAM”), something on which Getty Images had sought to rely in their Opening Skeleton. I determined that the existing pleading did not include reference to CSAM, a determination with which the Court of Appeal agreed ([2025] EWCA 749). On the third day of trial, Getty Images made an application on short notice to amend the Particulars of Claim so as expressly to include reference to CSAM. No attempt was made to appeal my decision to dismiss that application for reasons set out in my *ex tempore* judgment ([2025] EWHC 1450 (Ch)).

The Issues for Trial

75. Going into the trial, the court was provided with a List of Issues running to 57 different issues. By the close of trial, the List of Issues had been very substantially narrowed, albeit that it remained a somewhat cumbersome document which (as the parties accepted) did not really assist in elucidating the real points in issue in the case. With this in mind, and at my suggestion, the parties provided me with Decision Trees for both the Trade Mark Infringement Claim and the Secondary Infringement Claim designed to provide a route map for the court through the key issues. These have been most helpful and I am grateful to the parties for their cooperation in producing these documents. Nevertheless, I have kept an eye on the List of Issues with the aim of

ensuring that all outstanding issues between the parties have been addressed in this judgment, where appropriate.

Representation at Trial

76. Getty Images was represented at trial by Ms Lindsay Lane KC, Ms Jessie Bowhill and Mr Joshua Marshall, instructed by Fieldfisher LLP (“**Fieldfisher**”). Stability was represented by Mr Hugo Cuddigan KC, Mr Edward Cronan and Mr Henry Edwards, instructed by Bird & Bird LLP (“**Bird & Bird**”). I am extremely grateful to the teams on both sides for their helpful submissions and their cooperation throughout the trial in finding ways to assist the court in dealing with so many complex issues.

(D) THE WITNESSES AND EVIDENCE

77. The court heard from a total of 19 witnesses over the course of only a few days. They were cross-examined with concision and focus. Some of their evidence is of course no longer relevant to the issues that remain to be decided. With only one exception identified below, I formed the impression that all of the witnesses were honest and that they were all trying their best to assist the court.

Getty Images’ Witnesses of Fact

78. Getty Images called 10 witnesses:

- i) **Heather Cameron**, Director of Legal Risk Mitigation at Getty Images (Seattle) Inc, a management company which supports various entities in the Getty Images Corporate Group, including the First Claimant. Ms Cameron has worked at Getty Images for 26 years. She has provided four witness statements in these proceedings which (amongst other things) cover the history of Getty Images, the Getty Images Websites, the Getty Images Database and the investment made in it, the Getty Images and iStock watermarks; Getty Images’ AI Generator, damage to Getty Images’ business and Getty Images licensing practice. In her fourth statement, Ms Cameron gives evidence of an analysis of Copyright Works and Visual Assets identified in the LAION datasets, designed to identify the number of Copyright Works and Visual Assets on which Stable Diffusion was trained. Although Stability contends that no weight can be placed on the evidence in her fourth statement, it makes no personal criticism of Ms Cameron in relation to that evidence.
- ii) **Andrea Gagliano**, Senior Director of Artificial Intelligence and Machine Learning at Getty Images (Seattle) Inc. Ms Gagliano has worked at Getty Images since 2018. In her statement, Ms Gagliano addresses (amongst other things) her knowledge of, and investigation into, Stability and Stable Diffusion, including the datasets on which Stable Diffusion was trained and the means of accessing Stable Diffusion and generating an image from it. Ms Gagliano also gives evidence about Getty Images’ own AI model, GAI.
- iii) **David Stanley**, Vice President of Marketing at the Third Claimant, who joined Getty Images in July 2005. In his statement, Mr Stanley gives evidence about the ISTOCK and Getty Images logos and branding, his concerns

about Getty Images' reputation, and the potential confusion and damage caused by Stable Diffusion.

- iv) **Tanja Malnar**, Localization Project Manager for the Fifth Claimant. Ms Malnar gives evidence in her statement of the process by which contributors uploaded and licensed content to the Fifth Claimant through its website prior to 2017. Ms Malnar gave her evidence remotely from Canada.
- v) **Emma Varty**, a solicitor at Fieldfisher. In her first statement Ms Varty gives evidence as to the conduct of Getty Images' various experiments and the synthetic outputs generated. In her second statement Ms Varty explains the manner in which she identified the prompts used to generate Getty Images/ISTOCK watermarks in connection with the Getty Watermark Experiments. In her third statement, Ms Varty addresses the enrolment process that has been in place since 2017 for Getty Images' exclusively licensed contributors. Although Stability criticises the way in which Getty Images went about conducting the various experiments described by Ms Varty, Mr Cuddigan very properly did not criticise Ms Varty herself for the approach that had been taken.
- vi) **Mr Barwick**, a photographer and President of the Sixth Claimant and author of Sample Works E-K.
- vii) two freelance photographers: **Alberto Rodriguez** who captured SOCI Work **A8**; and **Alexander Livesey**, who captured SOCI Works **A5** and **A6**. Both gave their evidence remotely from the USA.
- viii) **Andreas Rentz**, a staff photographer for Getty Images DevCo Deutschland GmbH who gave evidence with the assistance of an interpreter. Mr Rentz captured SOCI Works **A10** and **A11**.
- ix) **Gregory Shamus**, a staff photographer employed by the First Claimant who captured SOCI Work **A7**. Mr Shamus gave his evidence remotely from the USA.

79. In addition, Getty Images rely upon statements from 3 witnesses who were not required to give evidence and whose evidence was therefore not challenged: **Matthew Prichard**, Senior Director of International Tax employed by Getty Images (Seattle), Inc; **Tessa O'Neill**, a Director of the Second Claimant since 2015; and **Jonathan Lockwood**, Vice President, Corporate Counsel at the Third Claimant. Given the narrowing of Getty Images' case at trial, very little of the evidence in these statements remains relevant.

Getty Images' CEA Notices

80. Getty Images have served four Civil Evidence Act Notices (the "**Getty CEA Notices**") identifying documents on which they intend to rely at trial:

- i) An updated CEA Notice dated 11 April 2025 attaching documents relied upon by Getty Images on the subjects of: (a) Getty Images' business; (b) copyright subsistence and ownership in connection with the SOCI Works and the Sample Works; (c) the Second Claimant and the issue of database right subsistence and

ownership; (d) Getty Images' Trade Marks; and (e) Stability and Stable Diffusion.

- ii) A CEA Notice dated 10 March 2025 attaching documents relating to the exclusively licensed works.
- iii) A CEA Notice dated 24 March 2025 attaching documents relevant to the SOCI Works.
- iv) A CEA Notice dated 27 May 2025 attaching screenshots showing the number of downloads of various versions of Stable Diffusion from Hugging Face.

Stability's Witnesses of Fact

81. Stability called 9 witnesses:

- i) **Cedric Wagrez**, Head of Data Strategy and Operations at Stability and an employee of Stability AI Japan KK. Mr Wagrez joined Stability in October 2023. In his first statement, Mr Wagrez explains his limited knowledge of the datasets used to train Stable Diffusion models released prior to his involvement with Stability. In his second statement he explains the investigations performed by Stability to establish whether certain images were used in training various versions of Stable Diffusion and he addresses the question of the number of Copyright Works used in training those versions of Stable Diffusion. Although Mr Wagrez's evidence was plainly honest (and Getty Images does not seek to suggest otherwise), I accept that the preparation of his first witness statement involved a very high level of generality which he then sought to remedy in his second statement.
- ii) **Daria Bakshandaeva** and **Mikhail Konstantinov**, research scientists who joined Stability as employees in the spring of 2022 and worked as part of what became known as the DeepFloyd Team. Both gave their evidence with the assistance of an interpreter. Although the importance of their evidence is now much reduced, I should observe that I agree with Getty Images that Mr Konstantinov was not a candid witness, that he was overly careful with his answers and that he resisted making appropriate concessions. This did not apply to Ms Bakshandaeva.
- iii) **Dennis Niedworok**, Head of Applied Research at Stability. Mr Niedworok joined Stability in December 2022 as a 'model fine-tuner' and became a member of the Applied team. In his statement he describes the process of "fine-tuning" a model and explains the work he was involved with at Stability.
- iv) **Kate Hodesdon**, a machine learning engineer who has been employed by Stability in the UK since February 2023. Ms Hodesdon works for the Inference Team, which is led by Anilkumar Bandari, Stability's Director of Engineering. In her statement Ms Hodesdon explains her work on inference, i.e. the process of running models once they have been trained.
- v) **Nojus Cebatoriunas**, an Information Technology specialist at Stability who joined as an employee in 2021 and is based in London. While no criticism is

made of his honesty by Getty Images, they contend (and Mr Cebatoriunas accepted in cross-examination) that the process he undertook to search for laptops available to Stability in the UK was quite obviously flawed. However, this is of little significance in the context of the remaining claims.

- vi) **Scott Trowbridge**, VP of Business Development who has been employed by Stability in the UK since March 2023. His role involves hiring individuals and supporting any sales-related or partnership opportunities. In his statement, Mr Trowbridge gives evidence about Stability's commercial customers and its licensing practices.
- vii) **Stephan Auerhahn**, was formerly a Principal Infrastructure Engineer in the API / Data / Infrastructure team at Stability based at all times in Los Angeles. Mr Auerhahn joined Stability in September 2022. In his statement he deals with data location and storage services, how inference works for the Developer Platform and DreamStudio, and the filtering used on those platforms. Mr Auerhahn gave his evidence remotely from the USA.
- viii) **Anilkumar Bandari**, Director of Engineering at Stability based in Seattle. Mr Bandari joined Stability in August 2023. In his statement, Mr Bandari gives evidence about the Engineering team. Mr Bandari gave his evidence remotely from the USA.

82. In addition, Stability relies upon the statements of 7 witnesses who were not required to give evidence and whose evidence was therefore not challenged. Five of these witnesses are existing employees of Stability: **Chantelle Palmer**, Head of People at Stability since June 2023; **Peter O'Donoghue**, Chief Financial Officer of Stability since the Summer of 2022; **Reshinth Adithyam**, a research scientist who joined Stability as an employee in December 2022; **Zachary Evans**, Head of Audio Research and an employee of Stability since June 2022; **Christian Burrows**, Head Security Engineer and an employee of Stability since March 2024. None of this evidence is of any significance in the context of Getty Images' narrowed case.

83. The remaining two witnesses whose evidence was not challenged by Getty Images are associates at Bird & Bird:

- i) **Aneesah Kabba-Kamara**, whose first statement explains: (a) how inference can be undertaken using Stable Diffusion v.1.4, v.2.1 and XL 1.0 running on a local machine; (b) how inference can be undertaken using Stable Diffusion v.1.6 via the Developer Platform; and (c) how inference can be undertaken using Stable Diffusion v.1.6 and XL 1.0 via DreamStudio. In her second short statement, Ms Kabba-Kamara addresses an omission from the information provided in her first statement.
- ii) **William Wortley**, who gives evidence (amongst other things) about the presence of certain images in the LAION-2B dataset which contain similar material to SOCI images A7, A10 and A11. This evidence is no longer of any relevance to the remaining issues.

84. Stability chose not to rely on a witness statement from **Oliver Toromanoff**, a former Lead Software Engineer who has now left his employment with Stability.

85. Until the day before he was due to give oral evidence, it appeared that Stability would also be relying upon the evidence of **Richard Vencu**, a former Stability employee who was involved with the training infrastructure at Stability between February 2022 and September 2024. One of Mr Vencu's two witness statements was relied upon by Stability's technical expert in his report and various references were made to Mr Vencu's evidence in Stability's Opening Skeleton.
86. However, Stability took the decision not to call Mr Vencu and so is no longer able to rely on the two statements signed by him in the proceedings (although Stability still relies upon documents that were exhibited to his statements, without objection from Getty Images). No reasons were provided to the court at the time for this decision. However during closing submissions Stability provided the court with a letter from Bird & Bird dated 27 June 2025 in which factors relevant to the decision not to rely upon Mr Vencu's evidence were explained.
87. In summary, these were that Mr Vencu's evidence related solely to Stability's computing resources, that Getty Images had dropped its claim in respect of "Final Training" of Stable Diffusion on Visual Assets in the UK prior to the trial and that his activities prior to working for Stability (when he had worked for LAION) were irrelevant to the matters before the Court. Furthermore, the letter relied upon the potential for Mr Vencu to be cross examined about an assertion (which Stability had itself made in cross-examination of Ms Gagliano) that the LAION database had originally included CSAM material which had been removed in RE-LAION (a more recent dataset compiled by LAION). Bird & Bird explained that the assessment required in order to determine whether it was fair and proper to ask Mr Vencu to attend for cross examination included consideration of whether Mr Vencu would need an opportunity to obtain his own legal advice prior to giving evidence.
88. Getty Images invites me to draw various adverse inferences by reason of Mr Vencu's absence from the trial, to which I shall return in due course.
89. In closing, Getty Images suggested that the court might draw adverse inferences by reason of the failure on the part of Stability to rely upon evidence from either Mr Mostaque or Tom Mason, Stability's Chief Technical Officer. However, in my judgment, this would be inappropriate. Aside from the fact that the relevance of their potential evidence was said by Getty Images in its opening submissions to go specifically to the Training and Development Claim (now abandoned), no attempt was made during the trial to establish whether these potential witnesses were alive, well and available to give evidence – an obviously relevant consideration (see *Efobi v Royal Mail Group Ltd* [2021] UKSC 33, per Lord Leggatt JSC at [41]). In closing, Getty Images did not seek to identify in any detail the relevant evidence that these missing witnesses would have been able to give, just as they did not seek to identify other existing evidence that might have a bearing on the points on which these missing witnesses might have given evidence. Accordingly I decline to draw any inferences by reason of the absence of these witnesses.

Stability's CEA Notices

90. Stability relies upon two CEA Notices (the "**Stability CEA Notices**"):

- i) a CEA Notice dated 23 January 2025, attaching a capture of the DiffusionDB Webpage. DiffusionDB is a dataset that (as at that date) contained 14 million images generated by Stable Diffusion using text to image prompts specified by real users. It was created by scraping user-generated images on the public Stable Diffusion Discord platform (a platform that allows users to generate images through an API); and
- ii) a CEA Notice dated 2 April 2025 attaching captures of the iStock Website, some documents from Companies House and a Photo Archive News Article.

The Technical Experts

91. Each side adduced evidence from eminent experts in the field of AI models (together “**the Experts**”):

- i) Getty Images rely upon the evidence of **Professor Hany Farid**, a Professor at the University of California, Berkeley with a joint appointment in Electrical Engineering & Computer Sciences and the School of Information. He is a member of the Berkeley Artificial Intelligence Research Lab and the Center for Innovation in Vision and Optics. He has a PhD from the University of Pennsylvania. He is co-founder of two companies dedicated (amongst other things) to developing software to detect manipulated images. His academic research focuses on digital forensics, forensic science, misinformation, image analysis and human perception. Over his 25 year academic career he has published over 200 academic papers in these areas of scholarship. He is the recipient of an Alfred P Sloan Fellowship, a John Simon Guggenheim Fellowship and is a Fellow of the National Academy of Inventors.
- ii) Stability relies upon the evidence of **Professor Thomas Brox**, a Professor for Pattern Recognition and Image Processing and Head of the Computer Vision Group at the University of Freiburg in Germany for the last 15 years. He has worked in the field of image processing and computer vision since 2002. He has an engineering doctorate from the Saarland University in Germany. He published some of the first works on deep learning that went beyond classification and he is known in the field for creating the convolutional encoder-decoder architecture known as the **U-Net**, which has become standard architecture for a very broad set of image processing tasks and is used in Stable Diffusion and other latent diffusion models. Professor Brox has worked as a part-time employee for AWS and is co-author of around 300 academic papers.

92. The Experts agreed a Joint Statement setting out the many areas on which they were able to agree. I asked at the PTR that they identify the remaining areas on which they disagreed and they then produced a further Joint Statement entitled “Statement of Areas of Disagreement”. In fact this latter document merely evidences the extent of the yet further agreement between the Experts. Nonetheless, both Experts were cross-examined and it became apparent during that cross-examination that there are areas of nuance in their respective reports on which they are not in complete agreement.

93. In reality, however, there are no conflicts of substance between the Experts in relation to any of the issues left in dispute. They both gave clear and helpful evidence to the court and neither side sought to (or needed to) criticise the opposing Expert.

94. In closing, Getty Images took issue with the instructions that had been provided to Professor Brox, submitting that they were in some respects deficient, a submission with which I agree. Specifically, Professor Brox had not been shown evidence which was capable of being relevant to the opinions he expressed in his report. In the event, however, this omission does not appear to me to undermine the evidence given by Professor Brox, whose responses under cross-examination were quite obviously direct, cogent and technically sound. As Stability submitted in closing (and consistent with his duties as an independent expert), Professor Brox readily agreed with propositions that were put to him as appropriate without regard to whether they might support or detract from Stability's case.
95. Shortly prior to the close of the trial and at my request, the Experts very helpfully produced an Agreed Technical Primer which provides a summary of the technology that is most relevant to the issues arising in this case. In so far as I have sought to summarise the technology behind Stable Diffusion in various parts of this judgment, I have drawn on the Agreed Technical Primer and am extremely grateful to the Experts for the additional work they have undertaken in putting it together.
96. Unless the contrary is stated, where I record the views of one or other expert in this judgment I understand those views to be unchallenged and I accept them.

Evidence of New York Law

97. Prior to the trial, the parties each served Statements of Case on New York law, relevant to the Licensing Issue, but neither side sought to rely on any expert evidence of New York law at trial. Instead, the parties have been able to agree upon a neutral statement of New York law, to which I shall return in due course.

The Trial Bundle

98. Although initially the subject of a dispute, it is now agreed between the parties that all disclosure documents included in the trial bundle are admissible as evidence of the truth of their contents pursuant to CPR PD 32 paragraph 27.2.

(E) LEGAL RESPONSIBILITY FOR STABLE DIFFUSION v1.X

99. Before I turn to look at the detail of the claims made by Getty Images in these proceedings, I must first determine a dispute between the parties over whether Stable Diffusion v1.x was released and/or made available by Stability via the CompVis GitHub and CompVis Hugging Face platforms. This dispute was captured by Issues 13-15 of the List of Issues which, broadly, focused respectively on (13) the development of v1.x; (14) the publication and making available of v1.x on these platforms; and (15) the question of where all versions of Stable Diffusion have been hosted.
100. As to the latter, it is accepted by Getty Images that the evidence at trial has established that DreamStudio and the Developer Platform are and were hosted outside the United Kingdom and that there is no clear evidence as to where the Github and Hugging Face platforms are hosted. Ms Lane confirmed in closing that there is now no need for the court to make findings on this issue.

101. Stability accepts (as I have recorded earlier in this judgment) that it made Stable Diffusion v1.x available to the public (for a limited time) via an API Platform and also that it made v1.4 available to the public via DreamStudio Beta – acts which are capable of giving rise to liability for trade mark infringement and passing off. However, I understand it to be common ground that access to the public via DreamStudio and the API Platform would not involve downloading the Model – an important consideration in the context of the Secondary Infringement Claim where Getty Images argues that it is the downloading of the Model in the UK that amounts to importation. Users would only download the Model (or more accurately the relevant Model Card, model weights and source code) if they gained access to it via the Hugging Face or GitHub platforms.
102. Thus issue 14 specifically concerns the question of whether Stability published and/or made available Stable Diffusion v1.1 to v1.4 on these platforms in August 2022, when v1.x was first released. Stability accepts that it provided computing resources to CompVis via its AWS Cluster and that it announced the release of Stable Diffusion on its website, but it says that each checkpoint of v1.x was made available by CompVis itself on the Compvis Hugging Face and CompVis GitHub pages, which were under the control of CompVis. Thus Stability denies that it is responsible for trade mark infringement/passing off and acts of secondary infringement of copyright (assuming they took place) in respect of *these* versions of the Model when accessed from *these* platforms. If Stability is right on this score, then it is accepted by Getty Images that it cannot succeed in its claim of secondary infringement of copyright in relation to v1.1 to v1.4 and it cannot succeed in relation to its claims of trade mark infringement and passing off in respect of any infringing watermarks generated using *these* platforms.
103. Save that Ms Gagliano confirmed in her oral evidence that v1.x was located on CompVis pages on the Hugging Face and GitHub platforms, there was no other witness evidence to assist on this issue. Getty Images relies primarily upon an admission from Stability in its Defence that it “launched” Stable Diffusion in August 2022, together with an analysis of the available contemporary documents, both as to the development and subsequent release of v1.x.
104. I need deal with the development of v1.x only briefly. Getty Images’ pleaded case is that Stability would not have given access to the AWS Cluster to Mr Rombach and/or CompVis for the purpose of developing an open source diffusion model (or any other AI model) “if the development was not understood to be done (wholly, jointly or in some other way)” on Stability’s behalf or for Stability’s benefit. Stability maintains that the investment was purely speculative and that “there is no indication that anything came in return for it other than goodwill”. It points out that Mr Rombach ultimately became a Stability employee and worked on v2.x.
105. Notwithstanding the lack of witness evidence in support of Stability’s factual position, Getty Images’ pleading alone certainly does not provide any real basis for a finding that Stability was itself directly responsible for the development of v1.x. Furthermore, the documents on which Getty Images rely in their written opening submissions provide no support for such a case. These include (i) Chat messages from May and June 2022 involving Mr Mostaque (identified as “**emad**”) and Mr Rombach together with researchers at CompVis and Stability employees; and (ii) a Chat message from February 2023 between employees of Stability. These Chat messages establish no more than that by June 2022 Stability was involved in discussions with Mr Rombach, Mr Esser (identified as “**pesser**”) and the CompVis team in connection with the training

and development of the Model; a proposition which appears to be accepted by Getty Images, which assert in their opening submissions that these documents show that Stability “was involved in and/or assisted” in the process of development prior to launch. The February 2023 Chat (which involves a discussion on what Stability’s employees viewed as misinformation contained in an article published in MIT Technology Review) takes matters no further; although it suggests that Stability did more than merely pay for use of the AWS Cluster in connection with the development of Stable Diffusion, it is inconclusive as to the extent of any direct involvement on the part of Stability.

106. In written closing submissions, Getty Images suggested that it might reasonably be inferred that the training datasets used to train Stable Diffusion v1.x are stored (in whatever form) within what have been referred to as “Stability’s S3 buckets” (“**the AWS S3 Bucket**”). While this might well be an available inference given that it is common ground that CompVis used Stability’s AWS computing resources to train v1.x and given that the available evidence suggests that Stability used the AWS S3 Bucket for the purposes of storage, I do not see that this elevates Stability’s involvement beyond mere assistance or involvement.
107. Absent a plea by Getty Images that Stability is jointly liable for any tortious acts of Mr Rombach and CompVis (and/or Runway) - a plea which was expressly disavowed in the recital to an Order of 28 March 2025 - mere “involvement” or “assistance” on the part of Stability in the development of the Model is plainly insufficient to give rise to any tortious liability (in the form of infringement of statutory intellectual property rights) that may arise by reason of its subsequent release to the public. As Stability submits, it is not enough for Stability to have facilitated the development of the Model by enabling access to (and providing the funding for) the AWS Cluster; it would also not be enough for Stability to have directed others to carry out any relevant act, or to have commissioned it, or even to have procured it – these are not acts which give rise to direct tortious liability in the case of statutory torts such as copyright or trade mark infringement (see *Lifestyle Equities CV v Ahmed* [2024] 2 WLR per Leggatt JSC at [92]-[95]).
108. In oral closing, I did not understand Getty Images to dispute Stability’s submissions on this point or seriously to maintain that issue 13 (development of the Model) was determinative of, or relevant to, any of the remaining causes of action.
109. Turning then to issue 14 (responsibility for release of v1.x), once again, Getty Images rely upon contemporaneous documents in support of their case that Stable Diffusion v1.x was released and/or made available for download via Hugging Face and GitHub as Stability’s own product, for which it was responsible and in respect of which it has direct liability. Specifically they rely upon the following documents in the period prior to 10 August 2022, when Stability announced the launch of Stable Diffusion on its website (emboldened words are emphasised by Getty Images in their submissions):
 - i) An email dated 21 June 2022 sent by Mr Mostaque to Chaotic Capital seeking support for “**our cutting edge open image model**”. The email also refers to “our ML team”, “the tech team” and says “we have dozens helping with the models”;

- ii) An email dated 28 June 2022 from Mr Mostaque to Amazon headed “Re: Stable Diffusion launch preparation<> Inferentia” providing a link to “the codebase for the Stable Diffusion Imagen variant **we are training**” at <https://github.com/pesser/stable-diffusion>”. The link appears to be to Mr Esser’s GitHub page. Later on the same day, Mr Mostaque sent a further email to Amazon in which he directed Amazon to another link saying: “[t]his one is a pretty simple variant of latent diffusion (CompVis team **part of team stability now**): <https://github.com/CompVis/latent-diffusion>”.
- iii) An email dated 19 July 2022 from Mr Mostaque to jmtrimble@gmail.com providing an update on Stable Diffusion and seeking further financial support. The email explains that “[w]e are nearing the first version of **our** stable diffusion image model for release”.
- iv) A Chat exchange dated 26 July 2022 between Mr Mostaque, employees of Stability and others in which he says “**We will** be announcing next Friday 5th August (yay) Stability AI” and “**We will** be releasing stable diffusion v1++ on compvis and **stability githubs** then”. Mr Mostaque goes on to say that “@pesser and @Robin Rombach would like to release as Apache”. Later in the same Chat, Christoph Schuhmann of LAION (“**Mr Schuhmann**”) (**spirit_from_germany**) asks “[w]hy not instantly v1+++512,...We still have v2++ and 1.5B to Release next ...And then imaginator”. Mr Mostaque responds a little later in the Chat “don’t trust model checkpoint proliferation ahead of launch...after launch all good”.
- v) A Chat exchange dated 8 August 2022 in which Stability employees are discussing the launch of Stable Diffusion. Mr Schuhmann says (in a series of consecutive messages): “so ‘Stable Diffusion’ will be announced as a cooperative Stability-EAI-LAION-Compvis effort, right?... @emad...maybe we should do a joint FAQ – video with a big influencer, that proactively answer the controversial questions...that could come up...a video dedicated to FAQs regarding the advent of Stable Diffusion and image generation in general”. Mr Mostaque replies “[y]es will send something lunchtime...Had to shift a bunch of stuff around...Don’t need an influencer...Can give regular updates and we have the dedicated community”.
- vi) Later in the 8 August 2022 Chat, Apolinário Passos, apparently an engineer at Hugging Face (“**Mr Passos**”) (**apolinariosteps**) makes the following observation:

“(Putting my PR/Branding hat) Seeing all the discussions in the Safety Squad server I am wondering can it become a problem that the academic name of the model and the name of the service/website/Stability add-ons are the same?

Like if someone takes the open model without safety safeguards, generates horrible things that would end up being filtered by a service API and posts saying ‘this is Stable Diffusion’ – if it was purely academic thing one could say ‘yeah, but here’s the model card with the biases acknowledgement and the indication to not

use in production etc' – like happened with the original Latent Diffusion

But if the ‘service’/‘branding’ of the inference service has the same name, if there’s a backlash to say ‘oh, but you used the Stable Diffusion open source model, not the Stable Diffusion service with the safeguards’ – may be hard to disentangle

Given that the ‘Stable Diffusion’ branding is already being established (sic) and getting traction as a service/a Stability branded thing, **I wonder if it would make sense that the model/academic codebase is released under a different name – and then there would be a mutually beneficial (for CompVis and Stability) separation between ‘Stable Diffusion’ – that can still be open source - ...and the ‘pure’, ‘vanilla’, academic Compvis model that maybe says it uses underlying but is not the same thing necessarily.”**

110. Getty Images contend that these documents show Stability taking “active steps” to raise funds commercially to exploit Stable Diffusion, treating Stable Diffusion as its own commercial product which it could launch and make available to the public as and when it saw fit, and making its own decisions about branding, packaging and marketing of the different versions of v1.x. This, say Getty Images, is indicative of an entity with control over and/or ultimate responsibility for the release of the product.
111. On a careful reading of these documents, I tend to disagree. Aside from the fact that Getty Images has no pleaded case of control or responsibility based on “active steps” in relation to branding, packaging, marketing or commercial exploitation, it is not at all clear to me that Getty Images’ interpretation of these documents is correct. I have already found on the evidence that Stability was involved in the development of Stable Diffusion together with Compvis and these documents are plainly supportive of that conclusion. But they say nothing determinative on the subject of whether Stability was in fact in control of, and responsible for, the release of Stable Diffusion v1.x on the CompVis Hugging Face and GitHub pages.
112. Certainly the documents establish a narrative of ongoing collaboration between Stability, Compvis and others – i.e. Mr Schuhmann’s express understanding that Stable Diffusion would be announced as a “cooperative Stability-EAI-LAION-Compvis effort” (which is not disputed by Mr Mostaque in his Chat reply, or by anyone else). The use of language referring to “our” Model and to actions that “we” will be taking is language that collaborators, working together, would use. Similarly the reference to the CompVis team now being “part of team Stability” appears to be an inclusive statement, designed to reflect the collaborative nature of the relationship between the two organisations; a relationship which is also reflected in a post made by Mr Mostaque nearly a year later, on 4 June 2023, in which he records that Stability was “a collaborator in the development of the first release of Stable Diffusion”. None of this evidence establishes on balance that Stability in fact took control over, and responsibility for, the launch of v1.x such that it will have direct liability for any tortious claims that may be established in respect of the release of v1.x to the GitHub and Hugging Face platforms, as Getty Images contend.

113. Mr Passos' intervention in the 8 August 2022 Chat expressly acknowledges the collaboration between Stability (as a commercial entity) and Compvis (as the academic research vehicle) and he certainly suggests the desirability of a separation between the two when it comes to branding. However, while it is unclear what knowledge he had as to the arrangements for the release of v1.x, his understanding appears to be that the Model/its academic codebase was a CompVis product.

114. On 10 August 2022, Stability posted an announcement in the following terms on its website (**emphasis added**):

“Stability AI and our collaborators are proud to announce the first stage of the release of Stable Diffusion to researchers. Our friends at Hugging Face will host the model weights once you get access. The code is available here [<https://github.com/CompVis/stable-diffusion>], and the model card is here [<https://huggingface.co/CompVis/stablediffusion>]. We are working together towards a public release soon. **This has been led by Patrick Esser from Runway and Robin Rombach from the Machine Vision & Learning research group at LMU Munich (formerly CompVis lab at Heidelberg University)** building on their prior work on Latent Diffusion Models at CVPR’22, combined with support from communities at Eleuther AI, LAION, and our own generative AI team.

The model builds upon the work of the team at CompVis and Runway in their widely used latent diffusion model combined with insights from the conditional diffusion models by our lead generative AI developer Katherine Crowson, Dall-E 2 by Open AI, Imagen by Google Brain, and many others. We are delighted that AI media generation is a cooperative field and hope it can continue this way to bring the gift of creativity to all.”

115. The announcement went on to set out quotes from various interested individuals, including the following:

“We are excited to see what will be built with the current models as well as to see what further works will be coming out of open, collaborative research efforts!” – Patrick [Esser] (Runway) and Robin [Rombach] (LMU)

“We're excited that state of the art text-to-image models are being built openly and we are happy to collaborate with CompVis and Stability.ai towards safely and ethically release the models to the public and help democratize ML capabilities with the whole community” – Apolinário, ML Art Engineer, Hugging Face”

“We are delighted to release the first in a series of benchmark open source Stable Diffusion models that will enable billions to be more creative, happy and communicative. This model builds

on the work of many excellent researchers and we look forward to the positive effect of this and similar models on society and science in the coming years as they are used by billions worldwide". - Emad, CEO, Stability AI

116. In a second announcement entitled "Stable Diffusion Public Release" which was published on the Stability website on 22 August 2022, Stability announced the public release of Stable Diffusion, providing a link to its own website at <https://stability.ai/stablediffusion> and again providing a link to the CompVis Hugging Face and CompVis GitHub pages. Later in the announcement, Stability directed potential users to DreamStudio "[f]or more control and rapid generation".
117. In my judgment, far from evidencing that Stability itself published the Model on the CompVis Hugging Face and CompVis GitHub pages, the 10 August 2022 announcement emphasises the collaboration involved in its development together with the leading role of CompVis and Runway. Importantly, the links to the code and to the Model Card which were provided in both the 10 and 22 August 2022 announcements are to *CompVis* pages on GitHub and Hugging Face; the *Stability* GitHub page that Mr Mostaque had apparently envisaged during the 26 July 2022 Chat was not in place at this time (although by 22 August 2022 Stability was providing access to v1.4 of the Model via DreamStudio (as it accepts)).
118. The CompVis GitHub page (URL <https://github.com/ComPVis/stable-diffusion>) is headed "CompVis/stable-diffusion" and a link is immediately provided to the Ommer Lab, the lab of Professor Ommer at <https://ommer-lab.com/research/latent-diffusion-models/>. Under the heading "Stable Diffusion" is the following post from CompVis:

"Stable Diffusion was made possible thanks to a collaboration with Stability AI and Runway and builds upon our previous work".

119. Later on the page (which includes detailed advice for downloading and sampling the Model together with details of the checkpoints that were being provided), CompVis explains that:

"Thanks to a generous compute donation from Stability AI and support from LAION we were able to train a Latent Diffusion Model on 512x512 images from a subset of the LAION-5B database"

"The weights are available via the CompVis organization at Hugging Face under a license which contains specific use-based restrictions to prevent misuse and harm as informed by the model card, but otherwise remains permissive. While commercial use is permitted under the terms of the license, we do not recommend using the provided weights for services or products without additional safety mechanisms and considerations, since there are known limitations and biases of the weights, and research on safe and ethical deployment of general text-to-image models is an ongoing effort. The weights are research artifacts and should be treated as such." (**emphasis added**).

120. The license referred to was a CreativeML Open RAIL-M license dated 22 August 2022 in which “Licensor” was defined as “the copyright owner or entity authorized by the copyright owner that is granting the license, including the persons or entities that may have rights to the Model and/or distributing the Model”. The copyright owners were identified as “Robin Rombach and Patrick Esser and contributors” (in contrast to the position in respect of v2.0 to which I have referred above). The CompVis Hugging Face page is a sub-page of the CompVis directory at <https://huggingface.co/CompVis> and is headed “CompVis/stable-diffusion-v1-1”.
121. Consistent with the evidence provided by the CompVis GitHub and Hugging Face pages is an article in The Verge by James Vincent dated 15 September 2022 attached by Getty Images to their CEA notice which refers to Mr Mostaque observing that the Model “is being released by a research institute as a generalized model”. The author of the article comments that by referring to a “research institute” Mr Mostaque “is referring to the fact that the technical license for Stable Diffusion has been released by the Ludwig Maximilian University of Munich’s CompVis lab, though Stability funded and shaped its development”.
122. In light of the available evidence, I reject Getty Images’ submission that the use of hyperlinks on Stability’s website directing end users to CompVis pages on GitHub and Hugging Face is evidence of Stability taking “active steps” to make available Stable Diffusion to the public in the UK via these pages. In my judgment the provision of these links was no more than the provision of information as to where a prospective user could find the weights, Model Card and code. It did not involve Stability actually making these resources available itself on these platforms. The fact that the code and Model Card were released on a CompVis page is entirely consistent with a collaboration between Stability and CompVis which involved Stability (as the commercial arm) funding the development of, being involved in, and publicising the launch of, Stable Diffusion v1.x, and CompVis (as the research and development arm) wishing to release the fruits of its labours in the form of the Model Card and source code on its own Hugging Face and GitHub pages from the outset. At this time, as the witness statement of Zachary Evans confirms, Mr Rombach was not an employee of Stability.
123. Further evidence which appears to me to be supportive of this conclusion is contained in an announcement published on the Stability website on 24 November 2022 (after Mr Rombach had become an employee of Stability) in which Stability announced the release of v2.0 and explained that:

“The dynamic team of Robin Rombach (Stability AI) and Patrick Esser (Runway ML) from the CompVis Group at LMU Munich, headed by Prof. Dr. Björn Ommer, led the original Stable Diffusion V1 release. They built on their prior work in the lab with Latent Diffusion Models and got critical support from LAION and Eleuther AI...Robin is now leading the effort with Katherine Crowson at Stability AI to create the next generation of media models with our broader team”.
124. On this occasion, (in contrast to the release of v1.x) Stability refers anyone interested in the release of v2.0 of the Model to “**our GitHub**: <https://github.com/Stability-AI/Stable> Diffusion” (**emphasis added**) and states that “[w]e are releasing these models

into the Stability AI API Platform (platform.stability.ai) and DreamStudio in the next few days". The Stability GitHub page refers to the fact that:

"The weights are available **via the StabilityAI organization at Hugging Face** under the CreativeML Open RAIL++-M License" (**emphasis added**).

125. On balance, having reviewed the available documentary evidence, I consider there to be no basis for a finding that Stability actively made v1.x available and/or published the source code and model weights for v1.x to the public via the CompVis Hugging Face and CompVis GitHub platforms, and/or that it was the "moving force" behind this publication, and I reject Getty Images' case to that effect. I do not consider the fact that Stability accepts that it collaborated in the development of v1.x, that it was involved in its launch, or that it published the Model elsewhere, to be inconsistent with this conclusion and, as I have already said, there is no case to the effect that Stability is a joint tortfeasor together with CompVis (or anyone else) for any allegedly infringing acts.
126. In closing, Getty Images sought to rely upon evidence from Mr Wagrez to the effect that when he was putting together his first witness statement he had understood that he was being asked about what data had been used to train v.1, v.2 and XL of the Model. It was put to him in cross examination that he therefore understood that Stability was "responsible for those models", to which he responded "Yes". However, I do not consider this evidence to be of any real assistance. Aside from the fact that Mr Wagrez did not join Stability until October 2023 and so can have no direct knowledge of the details of the release of Stable Diffusion v1.x, the question was posed in the abstract, without any attempt to explain what was meant by "responsible for those models". Given that Stability accepts that it did make v1.x available via access mechanisms other than the CompVis Hugging Face and Compvis GitHub pages, Mr Wagrez's evidence takes matters no further. Similarly, although it was suggested by Getty Images that Mr Auerhahn "did not distinguish" in his evidence between v1.x and later versions of Stable Diffusion and therefore the submission was made that "it is clear" that he thought that responsibility for Stable Diffusion resided with Stability, regardless of the version or sub-version, Mr Auerhahn's questioning was not directed at the question of whether Stability was "responsible" for releasing v1.x via any particular access mechanism.
127. I note in passing that the parties agreed upon an issue for disclosure specifically addressed to the question of Stability's responsibility for the publication of Stable Diffusion on CompVis GitHub and CompVis Hugging Face and identified various potentially relevant categories of documents, but, aside from the documents to which I have referred, no other evidence is relied upon by Getty Images. I reject Ms Lane's oral submission that the use of the CompVis page was mere "happenstance".
128. I also reject Getty Images' case (advanced in opening submissions) that Stability is responsible for publishing and/or making Stable Diffusion v1.x available via Mr Esser's Github page prior to the official launch of v1.x on the CompVis GitHub page. There is no support for such an allegation in the Particulars of Claim and it appears to be based upon the content of the 28 June 2022 email to Amazon referred to above. Getty Images provided no explanation as to why publication on Mr Esser's page could amount to publication to the public by Stability and I did not understand the point to be pursued in closing.

129. Finally I note that the question of responsibility for v1.x was an issue on which Getty Images invited me to draw an adverse inference by reason of the absence of Mr Vencu as a witness. Ms Lane drew my attention to *Efobi v Royal Mail Group Ltd* [2021] 1 WLR 3863 per Lord Leggatt JSC at [41], and submitted that she would have wanted to question Mr Vencu on the issue of the release of Stable Diffusion v1.x, not least because his name appears on various relevant documents. Ms Lane is right on that score – Mr Vencu was a party to various of the Chats to which I have referred above and occasionally contributed to them. Furthermore, as Getty Images point out, Mr Vencu was employed by Stability from February 2022 and his role involved providing infrastructure, compute and storage services to research teams. No doubt he would have had relevant evidence to give on the development and training of Stable Diffusion and it is certainly possible that he would have had material evidence to give on the issue of “making available to the public”.

130. Always subject to my being satisfied with the explanation for Mr Vencu’s absence, then I accept that if Getty Images had established a *prima facie* case to the effect that Stability released v1.x to the public via the CompVis Hugging Face and GitHub pages, then it may have been appropriate to infer that nothing Mr Vencu would have said in evidence (had he been called) would have undermined that case. However, in my judgment I need not consider the matter further. As set out above, the available evidence goes the other way and, in the circumstances, I do not consider it appropriate to draw any inference by reason of Mr Vencu’s absence and certainly not an inference that would be inconsistent with the evidence contained in the contemporaneous documents.

Conclusion on legal responsibility for v1.x

131. In conclusion, for the reasons set out above, I find that Stability bears no direct liability for any tortious acts alleged in these proceedings arising by reason of the release of Stable Diffusion v1.x via the CompVis GitHub and CompVis Hugging Face pages.

132. I note, however, that there is of course no dispute that Stability *did* make v1.x available via an API Platform site and (in the case of v1.4 only) via its own DreamStudio platform. All references to v1.x hereafter are intended to be references only to v1.x as made available via these two access mechanisms.

133. It is common ground that Stability is responsible for the release of the other versions of Stable Diffusion in dispute in these proceedings.

(F) THE TRADE MARK INFRINGEMENT CLAIM

Introduction

134. Getty Images allege that Stability has committed acts of trade mark infringement pursuant to sections 10(1), 10(2) and 10(3) of the TMA by reason of the fact that it has affixed watermarks (referred to hereafter as “**the Sign(s)**” or “**watermarks***”⁶), which are identical or similar to the Marks, to the synthetic image outputs of Stable Diffusion and used those Signs in relation to services, namely the provision of synthetic image

⁶ From this point on, I have adopted the approach of describing synthetically produced watermarks as “**watermarks***”.

outputs. Getty Images accept that not every synthetic image produced by Stable Diffusion produces a Sign which is similar or identical to the Marks, but they say that they have adduced evidence that some do.

135. The question of whether the use of a sign infringes a trade mark under these sections falls to be assessed as at the date that the use of the sign complained of was commenced (*Levi Strauss & Co v Casucci SpA* (C-145/05) [2006] E.C.R. I-3703 at [13]). It is impossible to identify when each Model will have generated an allegedly infringing watermark*, and so I accept Getty Images' case that the relevant dates for present purposes must assume infringement from the date of release of each Model. In other words, August 2022 for v1.x; November 2022 for v2.0; Mid 2023 for SD XL and November 2023 for v1.6. Although Stability complains that this approach is undisciplined in the context of identifying the use on which Getty Images relies for the purposes of infringement (because it relies upon use of the Models *per se*, rather than use of an infringing sign), nevertheless, I consider it to be the only sensible approach in the very unusual circumstances of this case. None other was suggested by Stability.
136. It is common ground that Getty Images must establish, on the balance of probabilities, that each version of the Model has generated at least one output with a watermark* containing a Getty Mark and a watermark* containing an ISTOCK Mark in the United Kingdom. It is important to differentiate between outputs produced by different versions of the Model because (as is also common ground) the Models have not all been trained on the same dataset and different filters have been applied to the training data. Thus while one version of the Model may generate synthetic images which include the Sign, another version may not.
137. There is no pleaded defence of "de minimis" from Stability and Mr Cuddigan confirmed in submissions that Stability did not intend to run such a case. Thus there is no scope for Stability to contend (as it appeared at times to do in its opening submissions) that detailed issues of infringement will only arise "if the court decides that the evidence of watermark* incidence 'in the wild', in the UK, is a problem of substance". However, in closing, Stability submitted that, in light of the evidence, a threshold issue arises, namely whether Getty Images have proved that any user of any version of the Model in the UK has ever been presented with a watermark* on a synthetic image generated by the Model. It is accordingly to this threshold question that I must turn first.

The Threshold Issue: Incidence of synthetically generated watermarks*

138. Stability admits in its Defence that "it may be possible" to generate synthetic image outputs which feature watermarks* and it accepted in closing that the Getty Watermark Experiments have shown that it is possible "to push" the Models in issue to generate watermarks*. However Stability says that these experiments were contrived and that when it comes to "non-legacy models" (which I understand to be a reference to SD XL and v1.6) the experiments failed. Stability relies upon a non-admission in its Defence as to the presence of watermarks* on output images generated by users of Stable Diffusion as the foundation for its argument (run at considerable length in closing) that Getty Images have been unable to prove at trial that any UK user of the various Models in issue has ever been presented with a watermark* on a synthetic image in the UK.
139. While Stability's primary case is that this broad submission applies to all versions of the Models, Mr Cuddigan accepted in closing that "there is a different position in

relation to different models” and that “there is a difference between v1.x and the others”. This difference arises because, as Stability acknowledged in its written opening submissions, “the incidence of watermarks* in outputs [from v.1x models] is non-trivial” and has “occurred to a non-trivial extent both in response to prompts designed by Getty to generate such watermarks*, and sporadically in relation to other prompts”. However, Mr Cuddigan explained in closing that this was intended as a reference to the incidence “worldwide” and that it involved no more than an acknowledgement that there is available evidence of commentary about the appearance of watermarks* in respect of v1.x, together with evidence that this has had an impact on users.

140. As for the remaining versions of the Model, Stability submits, in essence, that the available evidence is insufficient to satisfy Getty Images’ burden of proof, that much of the evidence of watermark* incidence has been produced using prompts that real world users would not use and that Getty Images has failed to advance any case based on the statistical probability of an infringing Sign being produced on a synthetic image generated from one or more versions of the Model.
141. Getty Images strongly resist this argument, pointing out that Stability’s Defence on the point is, at best, equivocal, and that Stability has failed to run any positive case asserting that watermarks* have not been generated by users of the various versions of Stable Diffusion in the UK. Getty Images submit that Stability has effectively already conceded the point on v1.x (notwithstanding Mr Cuddigan’s explanation in closing) and that, by reason of that concession, they had understood Stability to be advancing a case on wilful contrivance *only* in respect of v2.x, SD XL and v1.6 (and not v1.x).
142. Furthermore, Getty Images maintain that, in light of Stability’s non-admission in its Defence, Getty Images does not require much by way of evidence to discharge their burden of proof and they say that in fact they have ample evidence with which to do so. Owing to the absence of any plea by Stability that the generation of watermarks* by any of the Models in issue is *de minimis*, Getty Images submit that “it is certainly not necessary for Getty Images to prove that there is a particular probability of any of the Models generating watermarked images”.
143. To look at the threshold question in more detail, I must first examine the available evidence as to the generation of watermarks* generally by the Models. That evidence falls into four main categories: (i) the Expert evidence as to the scope for generation of watermarks*; (ii) Annex 8I to the PoC; (iii) the Getty Watermark Experiments and Annex 8H to the PoC; and (iv) Evidence of watermark* generation “in the wild”.

*(i) The Expert Evidence as to the scope for generation of watermarks**

144. In the Agreed Technical Primer, the Experts record that models such as Stable Diffusion can be prone to what is called **memorization** – in simple terms, the reproduction of an image used in training:

“The network’s weights are optimised on the training data, but its goal is to perform well on previously unseen data. In the context of Stable Diffusion, unseen data means new random noise patterns and/or new text inputs. To work reasonably on such new data, the network must be able to ‘generalise’: to

recognise and understand the general patterns and rules in the training data and be able to apply them in a different context.

If a network has been trained for too long on the same training data or an insufficiently diverse training data, it can be prone to '**overfitting**'. Overfitting occurs when the network uses its weights or part of its weights to memorize the individual training images rather than representing a large set of training images jointly with these weights. Overfitting is characterised by small errors on the training data, but a high error rate on new, unseen data. Overfitting is an undesired feature in machine learning, which engineers try to avoid.

Deep networks can both generalize and memorize at the same time. In such case, the network uses most of its weights to represent general patterns in the data, but uses some part of its weights to memorize individual patterns. The presumed primary cause for memorization is duplication of training data, either by explicit duplication or by training the network for too many **epochs**⁷, in conjunction with patterns that cannot be easily represented together with other patterns in the dataset – so-called “outliers””.

145. In their Joint Statement, the Experts record the extent of their (very considerable) agreement on the subject of watermarks*, which they confirm can be generated by the Model in response to certain prompts owing to the memorization process:

“We agree that Stable Diffusion does generate images with what appears to be a Getty watermark (albeit often distorted) and that this is due to the fact that the model was trained on some number of images containing this visible watermark.

...

We agree that the likelihood of a watermark appearing depends on at least the frequency with which the watermark appears in the training data and the user-specified prompt. We also agree that determining the precise likelihood is difficult because of the complexity of the influence of the prompt and also because of the sheer number of representative prompts and images generated per prompt that would need to be evaluated.

We also agree that certain prompts will generate a watermark with a high frequency, while other prompts are unlikely to generate a watermark.

With respect to how many images a model would need to see before it begins to reproduce a visible watermark, we note that

⁷ An epoch is a complete pass of all the data in a training dataset during the training process.

in the Carlini study⁸, they find memorization with duplication between 200 and 3000 duplicate images. While we cannot say for sure if this frequency would be the same for a watermark, it provides some insight into the frequency with which a model would need to see a watermark in the training data before it begins to reproduce it.

...

We note that in order for a watermark to be produced it is likely that the model needed to be trained on a diverse set of images/captions each containing a watermark. If it was only a single (duplicated) training image with a watermark, the model would memorize the whole image and not just the watermark”.

146. The Experts reiterated in their Second Joint Statement that “it is most likely that the dominant reason for memorization is data duplication during training either through duplication of the same image in the training dataset, or repeated exposures of one image during training”. In response to the subsequent question: “Are watermarks harder to generate than memorized images?”, the Experts responded:

“Whereas it takes multiple exposures to the exact same image to lead to memorization, memorizing a watermark likely requires multiple exposures to the same watermark regardless of the underlying image. It is not clear to us if one of these is easier or harder than the other. It seems to us that it is “easier” to find a prompt that shows a memorized image, because the image and its caption are reproduced in the training and so a caption with the appropriate keywords is more likely to generate the memorized image (assuming that we know the caption and/or keywords of the highly duplicated training image). On the other hand, it is less clear to us what prompt would generate a watermark since the exposure to the watermark would have been across many different prompts. In this regard, generating a watermark may be “harder””.

147. There does not appear to be any evidence of steps taken during the initial training and development of v1.1 of the Model to filter out watermarked images. In their first Joint Statement, the Experts agree that v1.2 of the Model “resumed training with a dataset that employed a watermark filter” (i.e. a filter designed to exclude the potential for the production of watermarks* in synthesised images by removing watermarked images from the training dataset). This is clear from the Model Card for v1.2, which records an estimated watermark probability of <0.5. Professor Brox accepted in cross examination that this means that “anything with more than a 50% chance of having a watermark has been excluded [from the dataset]”, although he explained that he was unable to say whether this probability threshold was conservative or not and how many watermarked images would remain after filtering. He accepted that the presence of the

⁸ The Carlini study (“Extracting Training Data from Diffusion Models”, Carlini et al. dated 30 January 2023) is relied upon by Getty Images and discussed by both Experts in their reports. The Carlini study showed that diffusion models memorize individual images from their training data and emit them when generating an output.

filter in the training of this version of the Model had plainly not removed the potential for watermarked* images to be generated and that he would expect outputs sometimes to generate watermarks*.

148. The Experts also agreed in their Joint Statement that, although not expressly stated in the Model Cards, it was likely that subsequent versions of the Model had employed a watermark filter. However, they also agreed that “even the best watermark filter will not be perfect and leave some images with watermarks in the dataset”. Professor Brox observed in cross examination that “[a]s surprising as it is, that we can build complicated models, it seems we still cannot build filters that are 100% correct”. The Experts expressed no opinion as to the likelihood of watermarks* being generated by a Model trained on a dataset which still retained some images with watermarks even after filtering. When asked about SD XL, Professor Brox said that he would assume that the filtering had got better, but he accepted that he did not know for sure, given that the Model Card contained no information as to any filtering that had been undertaken.

(ii) Annex 8I

149. Annex 8I to the Particulars of Claim provides evidence of various versions of Stable Diffusion producing synthetic images with watermarks* in real life. Stability accepts that Getty Images did not consent to the production of these images.

150. Annex 8I includes 26 images bearing Getty Images watermarks*. For many of these images the prompt and version of Stable Diffusion used to create the image is recorded as “unknown”. However, there are 8 images in Annex 8I in respect of which the full prompt is known. None of these prompts “corresponds” to any of the Marks, nor do any of them contain any of the Marks within them (and so they are all within the scope of Getty Images’ claim). There is no evidence that any of the prompts used corresponds to Getty Images captions, i.e. that real world users have in fact used Getty Images captions as text inputs to the Model.

151. Only 6 of the images in Annex 8I are identified as having been generated by a specific version of the Model – five are pleaded as having been generated by v2.1 (although Getty Images concede that one of these images, appearing at page 21 of Annex 8I is incapable of evidencing infringement owing to the extent of the blurring in the watermark*) and one is pleaded as having been generated by SD XL.

152. The four images on which Getty Images now rely in respect of v2.1 are at pages 7, 8, 24 and 25 of Annex 8I and show:

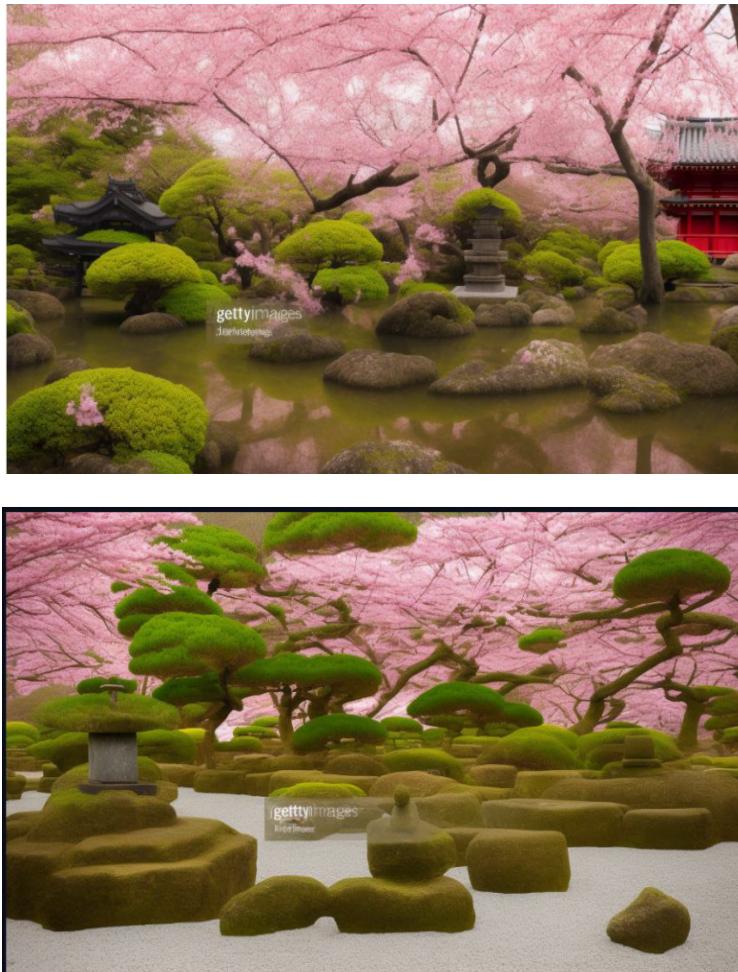
i) (in my words, because the prompt is unknown) a musician (“**the Musician Image**”) (page 7):



ii) (in the words of the prompt), “photo of two girls hugging” (“**the Girls Hugging Image**”) (page 8):



iii) (in the words of the prompt) a “japanese temple garden with blooming sakura and kami” (“**the Japanese Temple Garden Images**” (pages 24 and 25). I shall refer later in this judgment to the first of these images as “**the First Japanese Temple Garden Image**”:



153. I note that none of these three prompts is particularly complicated and none (as Professor Brox accepted in cross examination) has the particular features that he suggested in his report would increase the likelihood of a watermarked* image being generated. There is no reason to suppose that these are anything other than randomly picked prompts.
154. The Musician Image was posted on Reddit under the caption “getty images, really? Stable Diffusion 2.1” (“**the Reddit Musician Post**”). In the same Chat is confirmation from other users that they have experienced something similar: **Graucus**: “I can confirm. They have done this with mine”; **econopotamus**: “Just negative prompt ‘watermark’ or ‘text’ and it usually goes away in my experience”; **irateas**: “Yeah. This is common problem with 2.x+”; **tybiboune**: “they come up pretty often when you don’t specify any style”; **Grdosjek**: “You get those in 1.x too. Don’t be versionist :D”.
155. The Girls Hugging Image was also posted on Reddit under the caption: “I asked for a photo of two girls hugging and the AI tried to make a Getty images watermark. I didn’t ask for that but it is mildly interesting”.
156. The Japanese Temple Garden Images were attached to a Getty CEA Notice and appear to come from the Hugging Face (v2.1) platform. The user posting the images on 27 December 2022 (“**the Japanese Temple Garden Post**”) under the heading “Feedback on Training” says this:

“Just a thought...Many of the images being used to train stable diffusion come from getty images.

How do I know?

Prompt: Japanese temple garden with blooming sakura and kami

Half of the images generated have something that looks more than vaguely like the grey getty images stamp on it.

Even when you use the negative prompt “text”, some form of ghost proof stamp shows up.

Is there a way to train that out of the next version?”

157. Getty Images candidly accept that there is no evidence in Annex 8I of images with watermarks* having been generated in the UK, but it contends that, absent a case on the part of Stability that the Model behaves differently in the United States, or Australia, from the way in which it behaves in the UK, evidence that watermarks* have been generated by anglophone users is probative of the fact that they will also have been generated in the UK. In other words, as Ms Lane put it in closing, “evidence that it has happened in the real world somewhere is equally good evidence that it is likely to have happened in the UK”. Neither Expert sought to distinguish between use of the Models in different countries.
158. The one image in Annex 8I identified as having been generated by SD XL (“**the Steampunk Musician Image**”) was the product of a prompt in the following terms “A comic book style illustration of a steampunk musician who plays beautiful melodies using his steam powered instruments and devices, character design in style of ArtStation trending, DAZ, Octane render, steampunk, dynamic lighting, ((50mm cinematic film still)), 8k uhd, Intricately detailed, beautiful artwork by award winning comic book artists, She is holding a magical glowing lightbulb made of dna and swirling mysterious power crystals”:



159. Stability points out that the pleaded date of generation of the Steampunk Musician Image is 4 March 2023. However, as I recorded earlier in this judgment, Stable Diffusion XL Beta (the first version of SD XL) was not released to the public until 13 April 2023, over a month after the date of generation of this image. This image was shared in an internal Stability slack Chat on 4 March 2023 and appears to have been generated during the development of XL Beta. The relevant part of the Chat (discussing the Steampunk Musician Image) reads as follows:

“I know David is training better watermark predictor

It doesn’t seem prudent to launch DS with SD-XL and make a big deal out of a new model that produces Getty watermarks, right?

They’re currently dewatermarking.

I think this was a fluke, since they did remove Getty. Its things like people copying Getty images to Pinterest

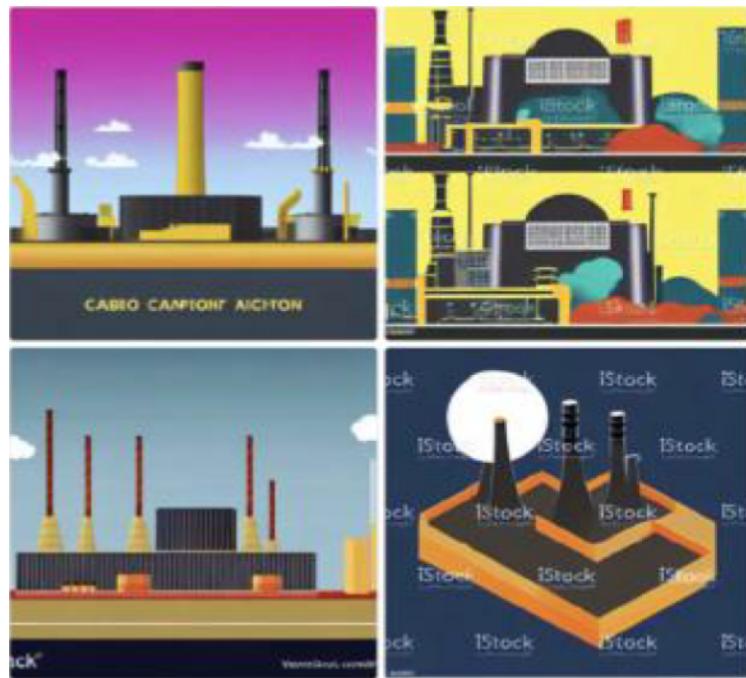
But got to be careful anyways

The smartest thing to do would be to revert to an earlier checkpoint and then resume training with the Dewatermarked LAION”

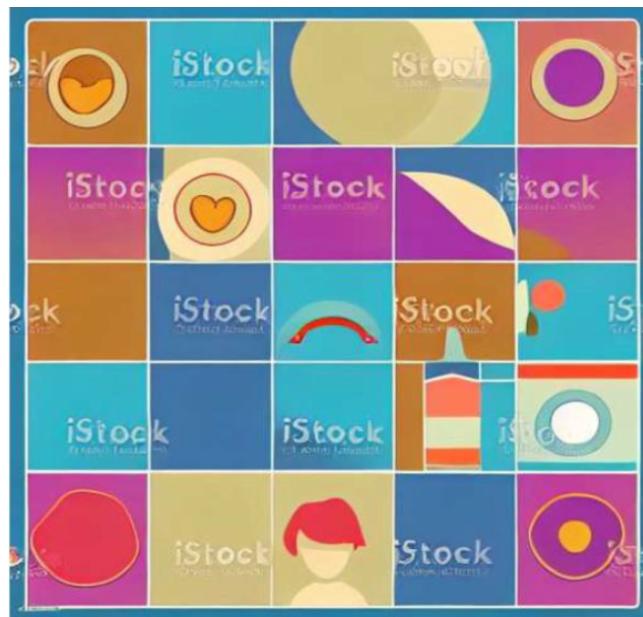
160. In light of this evidence, I find that the Steampunk Musician Image was created during the development of the SD XL Model such that the generation of a watermark* on that image cannot be attributed to the model weights which were subsequently released. It is clear that work was being done at the time the image was generated to “dewatermark” the dataset used in training and that these efforts were specifically directed at Getty Images watermarks. Accordingly, I do not consider the Steampunk Musician Image to provide any support for Getty Images’ case in relation to SD XL. This now appears

tacitly to be accepted by Getty Images. In closing, Ms Lane submitted that Annex 8I is evidence of the fact that “the early models of Stable Diffusion produced watermarks...”, thereby apparently accepting that Annex 8I provides no probative evidence in relation to either SD XL, or v1.6.

161. In opening, Getty Images sought to rely upon a table in which they made various, previously unpledged, assertions as to the likely version of the Model involved in generating the remaining images shown in Annex 8I in respect of which the version is currently identified in their pleaded case as “unknown”. They invited me to draw appropriate inferences. Stability objected to this on pleading grounds, but I do not need to address the pleading point in any detail. In closing, Ms Lane accepted that when one examined the inferences that the Court was being invited to draw, very few of them were really available owing to the fact that it remains impossible to be certain which precise version of the Model was involved in the production of the image. Thus:
 - i) proposed inferences that the Model used to generate the relevant image was likely to be “v1.x or v1.5” or “v1.x, v1.5 or v2.x” (Annex 8I pages 4 and 5) do not assist in circumstances where there is no issue in respect of v1.5 in these proceedings and it is impossible to rule out the possibility that the images were generated by v1.5. An inference that the Model might be either v1.x or v2.x is equally unhelpful where it is clear that the datasets on which these models were trained were different;
 - ii) proposed inferences that the Model used to generate the relevant image was likely to be a particular version “or anything prior” are also unhelpful for similar reasons (Annex 8I pages 6, 22 and 23);
 - iii) the proposed inference in relation to two of the images dated 10 March 2023 (pages 26 and 27 of Annex 8I) that they were likely generated by a “pre-release version of XL Beta” again takes matters no further because I am not concerned with a “pre-release version” of the Model.
162. That leaves (i) the images at pages 3 and 9 of Annex 8I dated 27 August 2022 and 28 August 2022 respectively; and (ii) the images generated on 1 October 2022 (at pages 10-20 of Annex 8I).
163. As to the four images at page 3 of Annex 8I generated on 27 August 2022, while I accept that it is impossible to say whether these were generated by v1.1, 1.2, 1.3 or 1.4 of Stable Diffusion, I consider it reasonable to infer from the date of the images that they were generated by a version of Stable Diffusion within v1.x (but not v1.5). Page 3 shows four generated images responding to the prompt “vector art with a flat material design style of factory with a carbon capture apparatus over the smokestack, solarpunk, optimistic” – they are set out below this paragraph. One of the four images very clearly includes the ISTOCK watermark*, while another also appears to do so, albeit less clearly. There is no evidence in relation to these images as to the mechanism used by the user to access the Model and no evidence as to where the user of the Model was when he or she created the images.



164. As for the images generated on 1 October 2022 (at pages 10-20 of Annex 8I) which are all conceptual images bearing the iStock watermark*, again I consider it reasonable to infer that these were generated by v1.x. The available evidence (in Stability's Responsive Statement of Case on Training and Development) suggests that v 1.5 was not released until after this date. I take the image at page 10 by way of example:



165. I am not prepared to draw a similar inference in relation to the image at page 9 of Annex 8I, created on 28 August 2022, which was attached to a message posted on Getty Images' SalesForce tool ("GI SalesForce") by a user who did not identify the Model that had generated the image, but simply referred to "AI systems like Stable Diffusion, DALL-E and Midjourney...". The image at page 9 of a woman in a group of other women bearing the Getty Images watermark* might have been generated by any one of these three different models; Stability is not responsible for DAL E or Midjourney.

166. Although I accept that Getty Images have never pleaded dates for the images at pages 3 and 10-20 of Annex 8I, I can see no real prejudice to Stability in drawing the inferences to which I refer above – particularly in circumstances where it concedes that the incidence of watermarks in v1.x is “non-trivial” and where the Getty Watermark Experiments at Annex 8H establish that v1.2, 1.3 and 1.4 generate iStock watermarks* when a prompt using the words “vector art” is used (as was the case for the image at page 3).
167. In conclusion, Annex 8I provides evidence that real users of v1.x have in fact generated synthetic images with iStock watermarks* (including by use of the prompt “vector art”) and that real users of v2.1 have in fact generated synthetic images bearing Getty Images watermarks* using apparently random prompts. While it is not clear that these users were based in the UK, I can see no difference in this case between users in the UK and elsewhere and none has been identified by Stability. There is no evidence to suggest that the Model works differently in different countries and Stability did not provide me with any cogent reasons why the evidence in Annex 8I should not be probative.

(iii) The Getty Watermark Experiments and Annex 8H

168. The Getty Watermark Experiments were undertaken by Professor Farid in conjunction with Fieldfisher. They involved text prompts being inputted into v1.2, 1.3, 1.4, 2.0, 2.1, XL 1.0 and v1.6 of the Model. The text prompts used were (i) verbatim prompts (i.e. broadly, prompts which were taken verbatim from captions on the Getty Images Websites which were chosen randomly by Ms Varty from captions used in a complaint filed by the First Claimant against Stability’s parent company in the United States); (ii) re-worded prompts (prompts generated by Professor Farid asking ChatGPT to re-word the original prompt); (iii) invented prompts (created by Ms Varty using the words “news photo” and “vector art” which she thought might produce watermarked* images); and (iv) prompts loosely inspired by other prompts (invented by Ms Varty and based on events that she was aware of, or were imagined by her).
169. The results of the Getty Watermark Experiments relied upon at Annex 8H, show outputs containing watermarks* being generated by:
 - i) V1.2: five images using verbatim prompts; nine images using re-worded prompts; and five images using “other prompts”. In respect of the latter, three of the images were generated using a prompt which included the words “news photo”, while one was generated using a prompt which included the words “vector art”. The image generated using the words “vector art” is the only image which includes an iStock watermark* as opposed to a Getty Images watermark*.
 - ii) V1.3: five images using verbatim prompts; six images using re-worded prompts and five images using “other prompts”. In respect of the latter, three of the images were generated using a prompt which included the words “news photo”, while one was generated using a prompt which included the words “vector art”. The image generated using the words “vector art” is the only image which includes an iStock watermark* as opposed to a Getty Images watermark*.
 - iii) V1.4: seven images using verbatim prompts; nine images using re-worded prompts; and six images using “other prompts”. In respect of the latter, three of the images were generated using a prompt which included the words “news

photo” and one was generated using a prompt which included the words “vector art”. The image generated using the words “vector art” is the only image which includes an iStock watermark* as opposed to a Getty Images watermark*.

- iv) V2.0: eight images using verbatim prompts and eight images using “other prompts”. In respect of the latter, two of the images were generated using a prompt that included the words “news photo”. All of the marks generated were Getty Images watermarks*. None of the images shows an iStock watermark*.
- v) V2.1: nine images using verbatim prompts; nine images using re-worded prompts and four images using “other prompts”. In respect of the latter, three of the images were generated using the words “news photo”, including an image responding to a prompt including reference to Miley Cyrus which is marked “Not Safe For Work”. All of the marks generated were Getty Images watermarks*. None of the images shows an iStock watermark*.

170. Although Professor Farid accepted in cross examination that attempts had been made to generate images bearing watermarks* using SD XL and v1.6, the Getty Watermark Experiments do not identify any synthetic images bearing watermarks* for these Models (notwithstanding that 2,600 efforts were made to do so using prompts designed for the purpose). Professor Farid also accepted that as one progressed through the Models “the frequency with which we saw watermarks was diminished over time”.

171. Nevertheless, a very small number of synthetic images bearing watermarks* were generated from v1.6 and SD XL by Getty Images during additional experiments that it undertook (as recorded in its NoE) in support of the Outputs Claim. Specifically:

- i) **SD XL 1.0:** one image of Donald Glover (included in Annex 8H) (“**the Donald Glover Image**”) bearing a distorted Getty watermark* (produced using a verbatim prompt);



- ii) **V1.6:** three images of The Gabba in Brisbane, Australia bearing Getty watermarks* (produced using a verbatim prompt, a re-worded prompt and a prompt which ends with the words “Getty Images stock photo”). This latter image does not appear to fall within the scope of Getty Images’ confirmation in the Order of 1 November 2023 that it would not rely upon prompts that contain

signs corresponding to the Marks) (“**the Gabba Images**”). By way of example, the image produced using the verbatim prompt (“**the First Gabba Image**”) appears below:



172. Professor Farid explained in his oral evidence that the purpose of the Getty Watermark Experiments was to understand when and if Stable Diffusion would generate images which included watermarks*. As he explains in his report, these experiments “are not capable of determining the precise probability of a user generating an output image similar to a training image or an output image with a visible watermark from the universe of possible output images”. He goes on to say that to determine precise probability he would need to generate a much larger number of images per prompt and run experiments across an enormous range of user prompts.
173. I did not understand there to be anything between the Experts on this point. Professor Brox described the Getty Watermark Experiments in his report as an “adversarial attack” and likened them to academic research (including in the Carlini study) in which the researchers had tried to identify whether a model is capable of being manipulated into generating specific outputs. In cross examination he explained that his use of the phrase “adversarial attack” was not a criticism and that the Getty Watermark Experiments were valid as “a proof of existence” – in other words an experiment to determine whether it is possible to find prompts that generate watermarks*. He rejected the suggestion that the Getty Watermark Experiments say anything about the probability of the Models generating images bearing watermarks*.
174. In their Second Joint Statement, the Experts agreed that “the Getty prompts likely overestimate the general prevalence of watermarks”. When asked about this in cross-examination, Professor Farid accepted that “it is perfectly reasonable to say the reason why we saw so many watermarks is that we were prompting the model with Getty captions and because we know that the Stable Diffusion models were trained on Getty assets, it was more likely to produce a watermark”. He accepted that he had insufficient data or information to estimate the extent of the influence of using Getty Images captions on the likelihood of generating a synthetic watermark*.

175. Professor Brox explained in his report that “[t]he very small number of synthetic images generated using Stable Diffusion 1.6 and XL which have been provided suggest that, even under the (likely) adversarial conditions of the [Getty] Watermark Experiment, these models rarely generate synthetic images bearing watermarks”. He expressed the view that he would expect that few watermarked images were contained in the training data, either “because a more effective watermark filter was used or because some other technique was used to identify and remove images containing Getty watermarks”. This view appears to be borne out by the available contemporaneous evidence in the form of the internal Stability Chat of 4 March 2023 to which I have already referred during which there is discussion of a “de-watermarking” process prior to the release of SD XL.

176. Turning to the arguments of the parties on the significance of the Getty Watermark Experiments, I understand it to be common ground that the outputs generated by the Getty Watermark Experiments included in Annex 8H are only of assistance in the context of the Trade Mark Infringement Claim if the types of prompts used to generate those outputs can be shown to be probative, or representative of real world use by real world users. That they establish a “possibility” of watermarks* being generated using particular types of prompts is not enough.

177. Stability contends that the Getty Watermark Experiments should be disregarded in their entirety on the threshold question, essentially because there is no evidence that Getty Images’ hand-picked, “contrived”, prompts have ever been employed by real-world users in the UK, that Getty Images has no case on the likelihood of watermarks* appearing (even when hand-picked prompts are used), that there are no experiments going to likelihood and no statistical analysis to support a case on probability and that most of the prompts represent an eccentric, out of scope use of the Models⁹ which is incompatible with the objectives of real world users. Stability’s pleaded case is that in normal use, users will seek to avoid generating images bearing watermarks* and that captions which correspond wholly or substantially to captions or alt-text for images from a Getty Images Website are inherently unlikely to be input by such a user by chance or during normal usage.

178. In support of this case, Stability relies primarily upon evidence contained in:

- i) a random prompt sample of 10,000 prompts submitted by real world users of Stable Diffusion XL 1.0 and v1.6 through the Stability Developer Platform API or DreamStudio on three dates: 20 March, 2 and 5 April 2025 (“**the Stability Prompt Sample**”). The protocol for this exercise was agreed by Getty Images which also selected the specific dates;
- ii) the first Stability Watermark Experiments addressed in its first NoE involving a selection of 1,000 text prompts from what is known as the “**Diffusion DB 2M dataset**”; namely a dataset of 2 million synthetic images and associated prompts from use of one of versions 1.1-1.4 of the Model. The randomly selected text prompts were input into v1.4, v2.1, XL 1.0 and v1.6 with 4000 inference requests being made in each case;

⁹ In other words, a use of the Models for which they were not designed, as is made clear in the Model Cards and elsewhere.

- iii) the second Stability Watermark Experiments addressed in its second NoE which used 525 verbatim text prompts taken from the prompts identified in AJG-10, an exhibit to Ms Gagliano's evidence ("AJG-10"), to generate 2100 synthetic images for each of v1.4, v2.1 and XL 1.0 and 2096 images for v1.6; and
- iv) its analysis of the (lack of) evidential value of the Annex 8H images.

179. Getty Images accept that they have no case on likelihood and that they have not attempted to conduct a statistical analysis, but as I have already indicated, they say there is no need to prove their case by reference to probabilities. They reject the suggestion that the prompts used for the Getty Watermark Experiments are contrived and assert that those prompts are representative of the type of prompts that a real world user might use, as well as being illustrative of the fact that any kind of prompt (of varying lengths) can generate a synthetic image with a watermark*. Getty Images assert that using verbatim (or substantially verbatim) prompts is something that a reasonable user of Stable Diffusion would do, because, for example, that user may wish to generate an image which is the same or similar to Getty Images' Content without paying a licence fee. They also contend that many prompts created by users without reference to Getty Images' Content are likely to correspond substantially to captions for such content "since these captions describe the Content, much of which is based on real places, people and events".

180. In addition to the expert evidence to which I have already referred and their own analysis of the Annex 8H images, Getty Images rely upon:

- i) a review carried out by Professor Farid of the Midjourney Discord Channel (a public online forum where users create and share the results of their image synthesis using an AI image generator called "Midjourney", a competitor of Stable Diffusion);
- ii) evidence of verbatim prompts input by users of GAI since the beginning of 2024 produced by Ms Gagliano at AJG-10 from an analysis of 470,000 prompts; and
- iii) their own analysis of the Stability Prompt Sample.

181. I begin by observing that (beyond Professor Brox's acknowledgement that any kind of prompt *may* generate a synthetic image with a watermark* and Professor Farid's evidence as to his review of the Midjourney Discord Channel, to which I shall come in a moment), the expert evidence to which I have already referred is of little real assistance to Getty Images on the question of the significance of the Getty Watermark Experiments. The Experts agree that these experiments show that it is possible for users of Stable Diffusion Models to produce watermarks*, but they also agree that the experiments say nothing about the likelihood of this happening. It would have been possible (as Professor Farid acknowledged) to run experiments (based on a large number of user prompts and images) designed to address the question of probability, but that has not been done.

182. The Experts' view that the Getty Watermark Experiments were designed only to establish "proof of existence" is supported by Ms Varty's evidence as to the means she employed in creating the Annex 8H prompts. In cross examination she confirmed that her only interest was in getting the Models to generate watermarks*. She was not

concerned with whether the synthetic images looked good or aesthetically interesting and if a particular prompt looked like it was generating absurd images that were ugly or unrepresentative or ridiculous, that would not have deterred her from continuing to prompt the Model with it if it was producing watermarks*. Ms Varty accepted that she had generated images for the purposes of the proceedings only and that (with the exception of the “vector art” prompts) all of her prompts were “out of scope” for use on the Models in the sense that the Model Cards explain that they were “not trained to be factual or true representations of people or events” and that accordingly using the Model to generate such images (as she had done for most of the images in Annex 8H) is out of scope for its abilities.

183. On the subject of the use made of Stable Diffusion by users, Ms Varty accepted that people may want images for a presentation or a project or an artwork and that they “want images which look good” and which are attractive and interesting. She also accepted that outputs which get human features comically wrong (as is the case with various of the outputs generated by Ms Varty) were likely to be a source of frustration and ridicule for people using the Models.
184. Notwithstanding that it is clear from Ms Varty’s evidence that there was no intention to reflect real world use in the Annex 8H prompts, the question remains whether that evidence is probative of the way in which real world users generate images using the Models such that it assists on the threshold question. To consider that question, I need to consider each of the different types of prompt used by Ms Varty with a view to determining whether, on balance, these prompts are probative or representative of real world use.

Verbatim Prompts:

185. Ms Varty chose these prompts from captions in the US proceedings as “an easily accessible public source of captions” which were believed to be in the LAION datasets used to train Stable Diffusion. They are all extremely lengthy and complicated, although, as Ms Cameron confirmed in her evidence, it is not unusual for Getty Images captions to be lengthy when they relate to editorial content (as all of the verbatim prompts do). One example which generated a Getty Images watermark* when using v1.2 is:

“U.S. President Barack Obama looks on during a review of military troops at a welcoming ceremony for French President Francois Hollande on the South Lawn at the White House on February 11, 2014 in Washington, DC. Hollande who arrived yesterday for a three day state visit, visited Thomas Jefferson's Monticello estate and will be the guest of honor for a state dinner tonight”

186. The image generated is set out below:



187. Would a real world user laboriously copy (or more realistically, perhaps, cut and paste) such a prompt into the Model in the hope of generating a similar (free, or much cheaper) image?
188. Professor Farid gave some evidence in his report that “[t]he experiments show that this reproduction [of watermarks*] happens without overly contrived or unexpected prompting by the user” and he went on to say that in his experience of working with AI models, and in the AI industry more broadly, “using captions from a stock image website is not an anomalous way to interact with an image generator”. It was by way of illustration of this proposition that he referred to his review of the Midjourney Discord Channel, observing that he had found “some examples of Getty Images captions used as prompts by users” and then providing an example in the form of an image of Ratan Tata, chairman of Tata Steel Ltd (“**the Tata Image**”).
189. Although Getty Images contended in opening that this evidence establishes users in the real world adopting verbatim prompts, Professor Farid’s candid responses to cross-examination on this point undermined that submission. Professor Farid very fairly accepted that the Tata Image was not an example of an image produced using a verbatim caption and that he thought he had not in fact found any examples of verbatim captions in his review of the Midjourney Discord Channel and that he could not say whether the caption had simply been re-written using ChatGPT. He also accepted that Midjourney is not a “particularly useful data point” and he explained that he could give no evidence on the frequency with which prompting of this type might be encountered.
190. I also consider that Getty Images’ reliance upon AJG-10 was seriously undermined during cross examination. AJG-10 runs to 15 pages of verbatim prompts entered by users of GAI (i.e. subscribers to the Getty Images service) – a total of 690 prompts (approximately 0.15% of the total prompts analysed). The exhibit lists the English word prompts entered by users since the beginning of 2024 “where five or more words match the words in an image caption associated with Getty Images’ content on Getty Images’ Websites”. Ms Gagliano confirmed in her evidence that the purpose of this exhibit was to demonstrate that there is a practice amongst users of GAI of entering Getty captions as prompts. She explained that a five word lower limit had been imposed because anything shorter than five words would not fairly have indicated that the text had been derived from a caption on the Getty Images Websites.

191. Stability accepts that AJG-10 supports the proposition that users of GAI do, on occasion, input Getty Images captions as prompts, and thereby also supports the general observation made by Professor Farid that using captions from a stock image website is not an anomalous way to interact with an image generator. However, it contends that AJG-10 provides no basis for an inference that users of Stable Diffusion behave in a similar way. On the contrary, Stability says that a comparison between AJG-10 and the Stability Prompt Sample illustrates that Stable Diffusion users behave differently from GAI users.
192. Stability's case was effectively accepted by Ms Gagliano in cross examination. I need not address this in any great detail as it was not addressed by Getty Images in closing (beyond a passing reference to Ms Gagliano's witness statement in Ms Lane's reply) and certainly Getty Images did not seek to gainsay the analysis that was put to Ms Gagliano during cross examination. Accordingly I can only assume that it is not in dispute.
193. For present purposes I record that Ms Gagliano accepted a scaling exercise that was put to her to the effect that (i) AJG-10 shows (for 10,000 prompts): 15 verbatim prompts of 5 words or more, 8 of 10 words or more, 5 of 15 words or more and 3 of 20 words or more; and (ii) by comparison, the Stability Prompt Sample data shows (per 10,000 prompts): 0 verbatim prompts with 10, 15 or 20 words or more and only 1 verbatim prompt with 5 words or more ("a flower in the middle of the desert"). Ms Gagliano also accepted that because the 470,000 prompts used to conduct the analysis for the purposes of AJG-10 covered all languages and that 690 prompts in AJG-10 are all in the English language, the figure of 0.15% referred to above underestimates the prevalence of English Getty Images captions amongst English GAI prompts such that on a like for like comparison the number of GAI verbatim prompts per 10,000 prompts would likely be higher. Ms Gagliano said nothing in re-examination to undermine these conclusions.
194. As for the one verbatim caption 'matched' prompt in the Stability Prompt Samples of 5 words or more, I accept Stability's submission that "A flower in the middle of the desert" is a trite phrase which is in no way unique or original to Getty Images. I also accept that it provides no proper basis for any inference of Getty Images as its source and thus does not support the proposition that real world users of Stable Diffusion use verbatim prompts copied from Getty Images Websites. I also did not understand this to be disputed by Getty Images.
195. Accordingly, as I was invited to do by Stability, I accept Ms Gagliano's evidence that these results support the proposition that Stable Diffusion users behave differently from GAI users and that, on the basis of the analysis that had been put to her, the proposition that her evidence was intended to support was "demonstrably wrong". It fails to assist Getty Images in establishing that verbatim prompts are in fact used "in the wild" by real world users of Stable Diffusion.
196. In her reply to Stability's closing submissions, Ms Lane submitted that the Stability Prompt Sample was too small to be of any assistance and that this point had been made by Fieldfisher before any testing was carried out in a letter dated 21 March 2025. I note however, that although the letter says that there will be a "limit to the probative value" of the Stability Prompt Sample and expresses Getty Images' concerns over whether it can properly be said to be representative of the "normal use" of Stable Diffusion, the

letter nevertheless agrees to it “in the interests of pragmatism and proportionality”. This is perhaps unsurprising given that Fieldfisher itself had proposed 5 sets of 2,000 prompts split into different word lengths in an earlier letter of 7 March 2025. Getty Images made no attempt to propose a larger sample size prior to the protocol being approved by the Court. In the circumstances, I accept Stability’s submission that Getty Images cannot now sensibly maintain that the criteria approved by the Court are not probative.

197. The Stability Watermark Experiments in respect of the Diffusion DB 2M dataset take matters little further on this issue. Stability asserts in its NoE that none of the 1,000 randomly selected prompts was a verbatim caption or a re-worded prompt. Getty Images declined to admit either allegation on the basis that “it is not reasonable, proportionate and/or practical” for them to do so. I am not prepared to find (as Stability invited me to do in opening) that on balance none of the 1,000 selected prompts was a verbatim caption or re-worded prompt. There is no evidence on which I could safely arrive at that conclusion and, as Getty Images correctly point out in their Reply NoE, it is unclear on what basis Stability has asserted that none of the selected prompts was a verbatim or re-worded prompt. No explanation is provided by Stability in its NoE as to any steps that it may have taken to establish the accuracy of this assertion and the mere fact that Getty Images have chosen not to conduct any reply experiments or any survey of the Diffusion DB database does not seem to me to be determinative. I note that Stability chose to say nothing about the Stability Watermark Experiments in connection with verbatim prompts in its closing submissions.
198. Insofar as the Stability Watermark Experiments illustrate that using a sample of 1,000 prompts generated from the Diffusion DB 2M dataset in respect of each of the Models v1.4, v2.1, SD XL 1.0 and v1.6 produced no images with watermarks*, I note and accept Professor Brox’s evidence that it was an attempt to estimate probability albeit that “the estimate is not super-precise because the number of samples was not too large. I agree it would be better if you used a million samples to do this”. This is entirely consistent with Professor Farid’s evidence in his report that determining a precise probability of reproducing a Getty Images watermark would require “much larger experiments” than those carried out by Stability. Professor Brox agreed that this experiment does not prove that real-world prompts do not produce watermarks* and that the images generated by the 1,000 prompts “likely under-estimate the general prevalence of watermarks”. He explained this in evidence on the grounds that “this Diffusion DB database may be a little biased towards artists”, a view taken also by Professor Farid in light of the content of the Model Card for the Diffusion DB 14M dataset on Github.
199. I accept the evidence of both Experts that the Stability Watermark Experiments are of extremely limited value in determining likely probabilities and certainly do not indicate that watermarks* will not be generated by real world prompts. They are of little assistance on the question of whether real world users in the UK have in fact generated watermarks* from any of the Models in issue in this case, just as they are of little assistance on the question of whether real world users in fact use verbatim prompts.
200. I can arrive at no determination as to the statistical probability of watermarks* being generated in any given situation based on any of the experiments undertaken by the parties. In closing, Stability sought to rely upon the second Stability Watermark Experiments whose intention appears to have been to establish that the AJG-10

verbatim prompts did not generate watermarks*. With the possible exception of one image discussed by Professor Farid in his report, the experiment did indeed draw a blank on watermarks* in respect of the Models covered, as Getty Images admitted in its Reply to Stability's NoE. However, the experiment establishes no more than that a small sample of verbatim caption prompts did not generate watermarks* for any of these Models. Professor Brox confirmed in his evidence that this has value as an independent sample from the distribution, while at the same time accepting that it would have been better to try a much higher number of prompts. Given that Stability accepts that the Models *can* generate watermarks* from various types of prompts, the only real value of the experiment is perhaps to highlight the absence of any more statistically significant experiment.

201. Tying the strands of this evidence together, beyond Professor Farid's general assertion in his report that the use of verbatim prompts is not "anomalous" in the context of interactions with an image generator (which finds no support in the example he then used of the Tata Image) I consider there to be no real evidence to support Getty Images' case that users of Stable Diffusion will copy Getty Images captions and paste these into the Model or, therefore, that users will have done this in the real world and generated watermarks*. Professor Farid's evidence finds support in AJG-10 in connection with the use of GAI, but the Stability Prompts Sample evidence shows that users of Stable Diffusion interact with it in a different way. I accept Stability's submissions that while there is evidence that the Stable Diffusion Models can be manipulated to produce watermarks* using verbatim prompts, there is no evidence whatever that this has happened in real life anywhere, including in the United Kingdom. None of the press articles, or third party materials, to which my attention has been drawn by Getty Images (and to which I shall come later) refers to users generating images using verbatim prompts and there is no evidence from even a single user in the UK that he or she has done so, much less that having used a verbatim prompt to conjure a synthetic image, a watermark* has in fact been generated.

Re-worded prompts

202. The re-worded prompts were generated by Professor Farid prompting ChatGPT with the prompt "provide a reworded version of this image caption: [relevant original Getty images caption]". Getty Images submit that the re-worded prompts are not contrived because they correspond substantially to captions for Getty Images' Content, which are often descriptive of the type of content (since it may be based on real places, people and events). In their opening submissions Getty Images asserted that "[i]t is likely that a user of Stable Diffusion would come up with something substantially similar to a caption in order to describe the type of content they wished to generate". However, there is simply no evidence whatever to support this assertion and certainly no attempt to establish a "likelihood" by reference to any statistical evidence. As I have already said, the evidence before the court is that these prompts were a deliberate choice by Ms Varty directly informed by her atypical use objectives.
203. I agree with Stability's submissions in closing that the case on re-worded prompts is necessarily dependent on the verbatim prompts case. This was supported by Professor Brox who explained in his oral evidence (which I accept) that he would not expect the re-wording by ChatGPT to make any real difference to the likelihood of watermarks* being generated because the concepts described in the text are what is important. I did not understand Getty Images to gainsay this evidence.

204. As I have said above, the verbatim prompts case posits that a user will visit a Getty Images Website, electronically copy a caption and then paste that caption into Stable Diffusion. As Stability points out, the re-worded prompts case involves an extra step: the Getty Images caption must be extracted, or at least read, in order for it to be re-phrased. If the approach adopted by Professor Farid were to be undertaken then the caption would have to be pasted into an AI model such as ChatGPT together with an appropriate prompt inviting the model to re-word the caption.

205. In closing, Getty Images put their case on verbatim prompts and re-worded prompts no higher than that “[a] user **may** copy and paste a caption from a Getty Images Website or amend a caption to vary what appears in the image...” (**emphasis added**). This much is not disputed by Stability. But there is simply no evidential support for the proposition that a real life user has in fact done this, or indeed that (having done so) watermarks* have been generated.

Prompts loosely inspired by other Prompts

206. The prompts at nos. 11 and 13-18 on page 25 of Getty Images’ NoE were invented by Ms Varty based on (i) events that she was aware of or were imagined by her; and (ii) featured people that she was aware of or were imagined by her, and which were loosely inspired by the verbatim captions she used elsewhere. Getty Images admit that, in one sense, these prompts are contrived because they were created by Ms Varty rather than arising naturally, but nevertheless they say that they are neither artificial nor unrealistic. They contend that these examples serve to demonstrate that (as the Experts accept) “with any kind of prompt, the user may obtain an image bearing a watermark*”.

207. The six “inspired prompts” used by Ms Varty are:

“11 Striker Harry Kane of Tottenham Hotspur scores equaliser during London Darby against the Chelsea at White Hart Lane on February 12, 2020 in Tottenham, London

13 Son Heung-min of Tottenham Hotspur celebrates after defeating Arsenal in North London Derby, on February 7, 2016 at White Hart Lane, London

14 Fred Smith of Brittingham Hotspur celebrates after defeating Arsenal in North London Derby, on 7 February, 2018 at White Hart Lane, London

15 Fred Smith of the Philadelphia Eagles celebrates after defeating the New England Patriots on February 4, 2018 in Minnesota

16 One Direction performs during the Super Bowl Halftime Show at Tulane Stadium on February 4, 2018 in New Orleans, Louisiana

17 Pop star Bruno Mars performs at Wembley Stadium, London on March 20, 2022 in, London

18 Pop star Dua Lipa performs at Wembley Stadium, London on March 20, 2022 in, London”

208. Stability accepts that the generation of images bearing watermarks* is possible, that real world users may generate such images and that such users may use very detailed prompts, but it contends that the Court has no means of determining that these, or equivalent prompts, have ever been deployed by any user of any version of the Models in the UK resulting in the generation of a watermark*.
209. It points out that, as Ms Varty accepted in cross examination, each of these prompts follows the grammar or structure that Ms Cameron confirmed in her evidence was the “house style” for Getty Images’ captions on news, sport, entertainment and contemporary Content. Stability submits that use of this structure would appear to be dependent upon knowledge of the Getty Images captions. I agree. In my judgment, reliance upon this category of prompts is dependent upon establishing that a real world user would wish to replicate the grammar and structure (together with the unconventional date format – at least for UK customers) of a Getty Images caption, while not copying it precisely and that, having done so, he or she would then generate a watermark*. There is simply no evidential basis for a finding that any real world user of Stable Diffusion in the UK has deployed this unusual syntactical arrangement or that, having done so, a watermark* has been generated.
210. Stability also points out that with the exception of the two ‘Fred Smith’ prompts (which are not concerned with real people) each of these prompts represents an “out-of-scope” use for the Models in that it is intended to produce a piece of misinformation: a fake depiction of an event (on a particular date at a particular location) that never took place. While it may of course be the case that “out-of-scope” use takes place in the real world, again there is no evidence to which I have been directed by Getty Images to show that “out-of-scope” use of this sort has in fact occurred (that the structure referred to above has been used) and that real users in the UK, or anywhere else, have generated watermarks* by such use. Certainly none of the images in Annex 8I (insofar as it is possible to see the prompts that have been used) appears to fall within this category. Further and in any event, as Stability points out, any user intending to produce fake imagery of this kind will of course be aware that he or she is engaging in fakery and will view the ensuing image with that aim in mind.

Invented Prompts

211. Ms Varty’s last category of prompts is “invented prompts”. Examples used for the purpose of the images in Annex 8H (at 19-24 on page 26 of Getty Images’ NoE) are:

“19 Miley Cyrus, performs onstage, during the 2008 American Music Award, news photo

20 Editorial news photo, Savoy School May 24, 2013 in Washington

21 Miley Cyrus at American Music Awards, news photo

22 News photo of MTV VMAs

23 Vector art spaceship

24 Obama, news photo”

212. These prompts are short and use two phrases which Ms Varty explains in her evidence she understood might produce watermarks*; namely “news photo” and “vector art”. Getty Images submit that the words “news photo” might be added by a user to a prompt “if they wanted a realistic image of a current event”, while the words “vector art” have in fact been used by real life users, as is clear from Annex 8I.
213. Stability accepts that the words “news photo” “might” be deployed by a user if he or she wanted a realistic image of a current event, albeit that it points out that this would be out-of-scope use of the Model. Further, it is clear from the examples in Annex 8H that the use of these words will, on occasions, result in the generation of a watermark*. Professor Brox accepted as much, explaining that the chances of this were likely increased because Getty Images Websites include many news photos.
214. Although Stability points out, correctly, that Getty Images has carried out no analysis of the frequency of use of this phraseology by real world users, or the circumstances in which it would, or would not, contribute to the appearance of a watermark*, there is one example of a real world user generating an image with a watermark* using this form of prompt in the evidence: an exhibit to Ms Varty’s statement shows two images of Donald Trump behind bars generated by Stable Diffusion v1.5 (which is not in issue in this case) with the caption “Old donald trump behind bars in a jail, news photo”. I reproduce one of those images below (“**the Donald Trump Image**”):



215. Ms Varty was unable to provide any information as to the identity of the user or why he or she had added the words “news photo”.
216. Although there is no evidence of any of the Models in issue being used in real life to generate images using a caption which includes the words “news photo”, I consider the

Donald Trump Image to provide evidence that a real world user of a diffusion model has in fact used the words “news photo” to generate an image. I do not find this very surprising. Professor Farid explained in his cross examination that it was very common for users of Midjourney Discord to “have stock phrases that they use over and over and over again on every caption because they have found that those particular words, like 85mm lens, F1.2 aperture, ultra realistic, give them the type of images they want”. I consider that it is reasonable to infer that for anyone wanting to achieve an “ultra-realistic” image, the words “news photo” might well be used in this way. Accordingly, I also consider it to be reasonable to infer that users of the Models in issue in the UK will also have used the phrase “news photo” to generate an image and, given the evidence in Annex 8H, combined with that of the Experts, I consider that, on balance, at least one real user of v1.2, 1.3, 1.4 and 2.1 in the UK will have generated an image using these words which includes a Getty Images watermark*. I take this view notwithstanding that some users may understand (from their reading of the Model Cards and/or other relevant material to which I shall return later) that the use of the words “news photo” amounts to “out-of-scope” use. All I really understand this to mean is that the Model was not designed to generate photo-realistic images of current events and real people; but I can see no reason whatever to suppose that this would prevent or preclude users from “playing around” with the Models to see what they are capable of producing (particularly given that the Model Cards highlight that the Models are intended for “research purposes”).

217. The words “vector art” (which are referring to a type of digital art that uses mathematical equations to represent images as a series of points, lines, curves and shapes) have been deployed by two real world users (as is clear from the images at pages 3 and 23 of Annex 8I). I have inferred that the first of these images was generated by v1.x. The second image is dated 28 January 2023 and so it is impossible to know for sure to which Model it relates and I have declined to draw an inference (although SD XL and v 1.6 can be ruled out owing to the date of the image). Ms Varty’s prompts using the words “vector art” produced synthetic images with iStock watermarks* for v1.2, v1.3 and v1.4 only. Getty Images has no examples in Annex 8H of synthetic images with watermarks* being produced by the use of the words “vector art” in respect of any other version of the Model. Accordingly, I consider that, on balance at least one real world user in the UK of v1.2, 1.3 and 1.4 will have generated an image using the words “vector art” which includes an iStock watermark*. There is no evidence on which I could make a similar finding in respect of any of the subsequent versions of the Model.
218. In conclusion:
 - i) The Getty Watermark Experiments and Annex 8H show that a user in the UK (i.e. Ms Varty):
 - a) of versions 1.2, 1.3 and 1.4, has in fact generated synthetic images with an iStock watermark* using prompts which include the words “vector art”; and
 - b) of versions 1.2, 1.3, 1.4, 2.0 and 2.1 has in fact generated a Getty Images watermark* using prompts which include the words “news photo”.

- ii) Given the evidence in Annex 8I, I consider the prompts in Annex 8H which use the words “vector art” to be representative of prompts which real world users would use. Given the evidence of real world usage of “news photo” to generate the Donald Trump Image (together with the Expert evidence referred to above), I consider the prompts in Annex 8H which use the words “news photo” to be representative of prompts that real world users would use.
- iii) The Getty Watermark Experiments provide no evidence as to the likelihood of Getty Images watermarks* and iStock watermarks* appearing in response to (respectively) “news photo” and “vector art” prompts and nor do they provide evidence as to the precise rendering of any watermark* that may appear in response to those prompts.
- iv) In so far as the Getty Watermark Experiments rely upon verbatim, re-worded and “inspired” prompts, there is no evidence of real world use of such prompts. Neither the GAI Experiment, nor the evidence of Professor Farid about his review of the Discord Channel provide any support for such use. There is also no support for this use in the additional available evidence of the generation of watermarks* “in the wild”, to which I shall turn in the next section of this judgment. Accordingly I reject the broad contention made by Getty Images in closing that the examples in Annex 8H should not be dismissed as “an isolated incident which only occurs through the wilful contrivance of litigation experiments”. There is no evidence on which I could do anything other than dismiss the majority of them (with the exception of the images produced using the prompts “news photo” and “vector art”) on this basis.
- v) The Getty Watermark Experiments and Annex 8H fail to produce any evidence whatever in respect of v1.6 and SD XL (including in relation to the use of either the “vector art” or “news photo” prompts). The only available evidence generated during the experiments in relation to these versions of the Model is to be found in evidence produced for the purposes of the Output Claim. However, that evidence was generated using verbatim and re-worded prompts. There is no evidence (including in the evidence to which I shall turn next of watermark* generation “in the wild”) that real world users of v1.6 and SD XL would use verbatim and re-worded prompts.

(iv) Evidence of watermark generation “in the wild”*

- 219. Annex 8I obviously contains evidence of real life generation of watermarks*. However, there is also some additional evidence to support Getty Images’ case.
- 220. I start with evidence internal to Stability which establishes that Stability was aware from the outset that there were problems with Models producing watermarked* images. In particular:
 - i) An internal Stability Chat dated 28 July 2022 in which Bill Cusick, former prompt engineer at Stability states: “There are still some watermarks showing up in archviz renders. Maybe once every 20-25 instances, so not often but when it does, it’s clearly revealing some kind of Getty or similar looking watermark”. This is pre-release of v1.x.

- ii) An email notification from Hugging Face to Mr Mostaque at Stability on 15 September 2022 with the subject “CompVis/stable-diffusion-v-1-4-original...Another training data quality issue – iStock watermarks”. The message says “Any prompt that includes the phrase ‘vector art’ will very frequently include a repeated ‘iStock’ watermark on the image. This probably isn’t the desired result and is likely due to an unauthorized scraping of iStock artwork into the training set”.
- iii) An email from a user dated 2 February 2023 to Team DreamStudio saying “it’s not the first time i generate images including the getty image watermark”. Team DreamStudio responds (“**the DreamStudio Email**”):

“Stable Diffusion (the models that DreamStudio runs on) was trained on a wide crawl of the internet and what that means is that occasionally generations will display a watermark due to the fact that watermarks were likely present in some of the dataset that the models were trained on.

While models are being trained, they are learning the features of the images that are in the dataset, so what can happen sometimes is that a certain set of words has a likelihood of displaying watermarks because it is thinking ‘I’ve seen a lot of images in this space, sometimes they have watermarks, so I should attempt to add a watermark to this image’.

This is an unfortunate byproduct that can come about by way of having such a large dataset and some of those images having watermarks in them that it learned from”.

- iv) An internal Stability Chat dated 4 February 2023 in which the participants discuss watermark removal and Tim Dockhorn comments: “When training w/o watermark conditioning, the model still generates samples with clearly visible watermarks. It’s in half of the data so I think that’s expected...”. This appears to be a discussion which is taking place during development and training, from which I infer that there was a recognition at that time that watermarked images made up about half of the dataset on which the next iteration of the Model was being trained.
- v) An internal Stability Chat dated 4 March 2023 (prior to the release of SD XL), to which I have already referred, in which the internal team discuss the current “de-watermarking” taking place in respect of SD XL and the fact that the watermark issue is a “blocker for launching”.
- vi) An internal Stability Chat dated 10 March 2023 (also at the time that training of SD XL appears to have been taking place) in which a comparison of SD XL ‘alpha’ and ‘Beta’ (both pre-release models) is taking place. Tom Mason says a decision needs to be taken as to whether to “switch in beta for the release”. Conner Ruhl, then a software engineer says “It’s specifically diagonals and the occasional Getty it seems to really be struggling with...”. Joe Penna responds “SDXL alpha (2.2.0) is making a watermark every other generation. I’ve yet to see one single watermark from ‘SDXL beta (2.2.2)’.

vii) A further internal Stability Chat from 10 March 2023 (again during the training of SD XL). Conner Ruhl says this: “Unless I’m specifically being targeted by some ghost in the machine...SDXL is not ready for prime-time or to be consider a blocker for launching...On a constant basis I’m seeing watermarks and furniture”. He then sets out links to examples, pointing out that these are all examples from the last day or two. He then comments “We are going to immediately be slammed by folks on the watermark issue alone”. Mohamed Diab later replies “just got 4 images all with watermarks, using one of the shuffle prompts and isometric style”. He also appears to provide links to examples. Conner Ruhl responds “Getty outta here”. Mr Auerhahn, who was involved in this chat, confirmed in his evidence that this was plainly a pun in recognition of the fact that Getty Images watermarks* were being generated. On the following day, Brian Fitzgerald observes that “results are more consistent than SDXL so far, no watermarks at all” – although it is unclear to what he is referring. Later in the same chat, Conner Ruhl says “Definitely would prefer versus the watermark/furniture model for the [DreamStudio] launch”. Later in the same chat, there is discussion around the fact that the problem with the model producing random furniture appears to be (as Scott Detweiler describes it) “somehow related to the ‘watermark killer’”, which I infer is a filter or additional functionality of some description, designed to remove the potential for the Model to generate watermarks. Scott Detweiler subsequently observes that “2.2 beta has so far not handed me random furniture”.

221. Pausing there – it is not always easy to understand exactly what is going on in these Chats. However, doing the best I can, I consider it reasonable to infer that v1.x had a tendency during training to generate watermarks*. This tendency was known to Stability prior to the release of v1.x and does not appear to have been ironed out on release, as the 15 September 2022 email notification (relating to iStock watermarks*) indicates. This is entirely consistent with my findings on the evidence in the previous section of this judgment.
222. Although attempts appear to have been made to filter the data that was being used to train v2.0 for watermarks (as confirmed by the Model Card), those attempts were not entirely successful – again, as my analysis of the evidence in the previous section also confirms. The DreamStudio Email appears to acknowledge as much when it talks about the “Models” on which DreamStudio was running at that time – these can only have been v1.x and v2.0.
223. It is also clear from this evidence that Stability encountered significant issues with the generation of watermarks in the development and training of SD XL. However, the evidence shows (and I find) that steps were taken to remove images bearing watermarks from the training data at this stage. Consistent with this is the fact that there is no internal Stability evidence indicating either that watermarks* continued to appear on synthetic images during real world use after the release of SDXL, or that there was a problem with later checkpoints of SDXL, or, indeed, v1.6. Accordingly, I infer from this evidence (together with the absence of any other evidence of watermarks* appearing “in the wild” in respect of these Models) that the inadvertent, uncontrived generation of watermarks* ceased to be a problem in real world use in respect of SD XL following its release and were never a problem in real world use in respect of v1.6.

I find that the only evidence of the generation of watermarks* in relation to these Models (discussed above) is contrived and unrepresentative.

224. In closing, Getty Images relied upon the evidence of Mr Auerhahn in cross examination to the effect that “[t]here were some issues with SDXL producing images that appeared to be watermarked” and that he was “possibly” aware of similar issues for “v1 and v2”. Mr Auerhahn accepted that the problem was “too prevalent” with SD XL. Bearing in mind that Mr Auerhahn was a party to various of the Chats to which I have referred above, it is not surprising that he recalled an issue with SD XL. However, his cross examination did not clarify when these concerns arose and, specifically, whether they arose pre- or post-release of SD XL. Mr Auerhahn was taken only to pre-release Chats during his cross examination. This is important because there is no evidence whatever of any watermarked* images being produced “in the wild” post-release and it was not specifically put to Mr Auerhahn that there were ongoing difficulties post-release of SD XL. He was also not asked about the evidence in the Chats which suggests that work was done to filter out images bearing watermarks from the training data and thus to preclude the possibility of the Model producing watermarks*.

225. Given that Mr Auerhahn was not asked in cross examination to clarify *when* he understood the problems to have arisen, I can see no basis to find that he was referring to the post-release period of SD XL when he identified the existence of problems with the generation of watermarks*. Much more likely, it seems to me, given the contemporaneous documentary evidence, is that he was referring to the problems that were being encountered pre-release of SD XL, and so I find. That would be entirely consistent with the fact that (i) Getty Images’ “proof of existence” testing was unable to identify a single watermarked* image for SD XL and v 1.6; and (ii) its Output Claim testing identified only the Donald Glover Image for SD XL and the Gabba Images for v 1.6 – albeit generated (as I have already found) by unrealistic and eccentric use that is unrepresentative of real world use.

226. Finally, there remain two additional sources of evidence on which Getty Images rely: (i) direct communications between members of the public and Getty Images on GI SalesForce; and (ii) Exchanges between members of the public on Reddit, a social news and discussion website (“**the Reddit Exchanges**”).

227. Getty Images have chosen to redact the names and contact details of all of the individuals who got in touch with its Salesforce team. They have not chosen to call any of these individuals or to permit Stability to view their details so that it might make a decision as to whether to call them. To my mind (beyond establishing that watermarks* have been generated by real world users) this limits the usefulness of the evidence in the sense that it is impossible to resolve any ambiguities. Instead I must do the best I can on the documents alone. There are three communications, all posted to GI SalesForce, on which Getty Images rely:

- i) A message from a user in the USA dated 28 August 2022 (“**the August 2022 SalesForce Message**”) attaching the image at 8I, page 9 to which I have already referred, saying:

“Is your leadership team aware that companies behind AI-generated imagery have used your database of imagery to help train their AI systems to generate photographs in order to replace

your business model? AI systems like Stable Diffusion, DAL E and Midjourney are using snapshots of the entire internet without the consent of copyright holders. Here is an image of an in-progress AI generated render that shows clearly that the AI was using your stock photo image" [an image with a watermark* is attached].

ii) A message from a user in the USA dated 1 October 2022 ("the **October 2022 SalesForce Message**") attaching various images (included in Annex 8I) on which the iStock image appears:

"I've been using Stable Diffusion and recently I came across a prompt that results in well over 50% of the images showing an iStock water mark, are these images legal to use in my projects? How does Stable Diffusion go about licensing material from you?"

iii) A message from a user in Latin-America dated 6 March 2023 ("the **March 2023 SalesForce Message**"): "I'd like to report a case of stable diffusion using Getty Images". The message then provides the prompt used and attaches an image with a watermark, albeit that the image has been generated by Stable Diffusion Online, which does not appear to have anything to do with Stability. Furthermore, this does not appear to be an anglophone user. I can only assume that it is for these reasons that the image is not included in Annex 8I.

228. These messages take matters little further than Annex 8I. Two of them are messages which attached the Annex 8I images. It is unclear whether the first message relates to Stability and I have already refused to draw an inference as to the Model involved for that reason. The image referred to in the third message certainly does not appear to have anything to do with Stability. I have drawn an inference in relation to the images attached to the October 2022 SalesForce Message and I note for present purposes that the images attached to that message were not the only examples that the user had been able to generate with watermarks*.

229. Getty Images rely only upon a few exchanges between members of the public on Reddit, two of which feed into the images at Annex 8I and two of which do not appear clearly to relate to images generated using platforms connected with Stability (albeit they evidence that the generation of watermarks on synthetic images in general was being discussed by users):

i) A post by Antique_Plane_130 from around August 2024 which asks "What's the best way to remove a watermark from an image?". However, as Ms Cameron accepted, the relevant post does not mention a Getty Images or iStock watermark. It refers to a "blurry watermark in one of the corners" and says that the user is using "automatic 1111". It is not clear that this is connected to Stability and there is no accompanying image. There is no reference to a Stable Diffusion Model and no indication that the user is based in the UK. I do not regard this post to be of any assistance.

ii) A post by NealJMD from 27 August 2022 which shows four images (used in Annex 8I at page 3 and illustrated earlier in this judgment at paragraph 163) and

says “Looks like Stable Diffusion was trained on watermarked images – when asked for vector art, it put the iStockPhoto watermark all over it”. The prompt is identified lower down in the Chat. I have inferred earlier in this judgment that this is a reference to v1.x, but there is no indication that the user is based in the UK. Other posts on this Chat include (from Musicguy1982) “I’ve gotten a few results with the Getty Images watermark although unreadable”; (from BunniLemon) “I tried the same prompt, and noticed that when I added ‘trending on ArtStation’ only one image showed up with the iStock’ logo...”; and (from namesareunavailable) “lately I experience awful lots of watermarks, even if I put watermark in the negative prompts. To me it has become quite annoying actually”.

- iii) A post (from Rare_Negotiation_544) on 31 October 2023: “Stable Diffusion generated a Getty Images watermark”. The embedded image is at page 6 of Annex 8I (and appears later in this judgment at paragraph 386) and the watermark* is very distorted. The model is unknown and there is no indication as to the location of the user.
- iv) A post (from lonewolfmcquaid) on 23 February 2023 asking why Stable Diffusion generates “accurate watermarks”, together with an embedded image showing a model and a badly distorted watermark. However, later in the exchange lonewolfmcquaid says that he made it with “playground” – a platform that is unconnected to Stability. In the same Chat, redpandabear77 says “I’ve made a ridiculous amount of images and I’ve never seen anything close to this. But I also don’t use the default models and I never try to make anything that resembles stock images. I just have a hard time believing that somebody got it looking that clear without trying to prompt for it”. Bobi2393 responds by sending a thread from this subreddit six months and three months earlier about other watermarks and other people confirming a similar result. Lonewolfmcquaid comments “i mean i was shocked as hell, it almost felt like winning a lottery cause what are the actual fucking chances lol”. This Chat again does not take matters much further on the threshold issue given that the image was generated on a platform that is unconnected to Stability.

230. There are a few other Reddit posts which produced various of the remaining images in Annex 8I, but again do not advance the position on the evidence any further.

231. In closing, Getty Images also relied upon an article by Stephen Jukic dated 18 January 2023 reporting on the commencement of litigation by Getty Images against Stability which includes an image with a distorted Getty Images watermark*. However, there is no indication which version of the Model was used to generate the image.

Conclusion on the Threshold Issue

232. Having regard to all of the available evidence considered above, I now turn to look specifically at each of the relevant versions of the Model to see whether the threshold question is satisfied.

Model v1.x

233. In support of its submissions in closing that Getty Images are unable to prove their case on incidence of watermarks* in respect of v1.x, Stability makes three points:

- i) First, that Stability has only been involved in a small proportion of the synthetic outputs generated by these Models. On Stability's case, which I have accepted, this court is concerned only with images generated by v1.x using the access mechanisms provided by the Developer Platform and DreamStudio (the latter for v1.4 only). Stability is not responsible for images produced by users accessing v1.x via the CompVis Hugging Face and CompVis GitHub platforms;
- ii) Second, Stability points to evidence in the form of an online article from Everypixel Journal dated December 2023 produced by Getty Images under a Getty CEA Notice which estimates that there are "as many as tens of thousands of Stable Diffusion-based models uploaded by [the] users" of GitHub, Hugging Face and Civitai – these Stable Diffusion-based models are not the responsibility of Stability;
- iii) Third, Stability relies on a further point in the same article to the effect that users of Stability's official channels generate 2 million images on a daily basis and that in over a year since its release, the number of images generated has reached 690 million out of a total potential number of 12.59 billion. Stability says that is less than 5.5% of the total images generated. Mr Cuddigan points out that users in the UK will be "a fraction of that", although there is no evidence to assist on what that fraction might be.

234. In light of the available evidence, I consider that, on balance, at least one user of v1.2, 1.3 and 1.4 accessed via the Developer Platform/API and (in the case of v1.4) also via DreamStudio in the UK will have come across a Getty Images watermark* and an iStock watermark* generated on synthetic images (*a fortiori* I consider that at least one user accessing via CompVis Hugging Face and CompVis GitHub will also have done so, should it be relevant to say so). While I cannot begin to identify how many times this might have taken place (and certainly do not have the evidence which would enable me to do so), I consider the weight of the evidence which I have set out in detail above to support such a finding.

235. I also consider it significant that Stability conceded in opening that the incidence of watermarks* in respect of v1.x is "non-trivial" and that it has occurred "to a non-trivial extent both in response to prompts designed by Getty to generate such watermarks and sporadically in relation to other prompts". Stability did not seek to withdraw its concession as to the "non-trivial" incidence of watermarks in v1.x in closing, but expressly maintained its position. While users of DreamStudio (for v1.4) and the Developer Platform (for v1.x) in the UK will only account for a small fraction of the total images generated across the world using Stable Diffusion, that small fraction is still likely to be a very considerable number of images and on balance I am satisfied on the available evidence (including the concession) that watermarks* will have been generated by such users. I note other relevant concessions from Stability during the trial as follows: "there is no doubt that there are plenty of examples of watermarks 'in the wild' amongst the billion or more images produced by versions 1.1-1.4 of Stable Diffusion"; that "we do accept there is some evidence of watermarks in the wild for v1

to v1.4”; and “[Ms Lane] says Stability was aware there was an issue with watermarks and we agree. It was something which happened”.

236. In respect of the generation of iStock watermarks* by v1.x (and in addition to the evidence of the Experts), I consider the following evidence to be particularly persuasive (i) the “real life” “vector art” prompt at page 3 of Annex 1; (ii) the “vector art” prompts in Annex 8H illustrating the generation of watermarks* in the UK for versions 1.2, 1.3 and 1.4; (iii) observations by users in the Reddit Chat of 27 August 2022; (iv) the 15 September 2022 Hugging Face email to Mr Mostaque and (v) the October 2022 SalesForce Message.
237. In respect of the generation of the Getty Images watermarks* by v1.x (and in addition to the evidence of the Experts), I rely upon (i) the “real life” Donald Trump Image generated using the words “news photo” together with an article in The Verge by James Vincent of 15 September 2022 which used this real life image observing that “[t]he presence of copyrighted imagery in Stable Diffusion’s training data is obvious from the program’s tendency to reproduce the ‘Getty Images’ watermark in certain pictures”; (ii) the “news photo” prompts in Annex 8H illustrating the generation of watermarks* in the UK for versions 1.2, 1.3 and 1.4 and (iii) the DreamStudio Email.

Model v2.x

238. I also consider that on balance at least one user of each of Stability’s v2.x Models accessed via DreamStudio and/or the Developer Platform and/or Hugging Face and GitHub in the UK will have come across a Getty Images watermark* generated on a synthetic image. Again, I consider the weight of the evidence which I have set out in detail above to support such a finding. Furthermore, I note Stability’s concession in opening (inevitable given the content of Annex 8I) that “it does appear there are some instances ‘in the wild’” and that “it is possible for users of v2.1 to generate images with watermarks* without contrivance” although Stability submits that the incidence of this is “insignificant” and “does not justify the court’s interference”. Given that there is no *de minimis* case pleaded by Stability, this concession is again important and amply justifies my finding that Getty Images have satisfied their burden in relation to v2.x. Although this concession was not repeated in Stability’s closing submissions, I did not understand it to be withdrawn and, in any event, in my judgment it correctly reflects the available evidence.
239. For v2.0 and 2.1 (in addition to the evidence of the Experts), I rely upon the following evidence: (i) the examples of real world generation of watermarks* in Annex 8I; (ii) the Japanese Temple Garden post; (iii) the observations made in respect of the Reddit Musician Post; (iv) the DreamStudio Email; (v) the Annex 8H “news photo” prompts illustrating the generation of watermarks* in the UK for versions 2.0 and 2.1 (combined with the Donald Trump Image which shows that “news photo” is a prompt that users use in real life – albeit on this occasion in relation to v1.5).
240. I am not swayed by Stability’s various experiments to take a different view on v2.x. The Diffusion DB experiment used only 1000 prompts and the AJG-10 experiment used only 525. While these experiments produced no watermarked images (with the exception of one image identified by Professor Farid in his report which shows two very distorted watermarks* in its bottom right-hand corner), they are not evidence that

v2.x does not in fact produce watermarked* images in the real world for all the reasons I have given above.

241. There is no evidence that iStock watermarks* have ever been generated in real life using v2.0 or 2.1 and accordingly I find that the threshold issue is not satisfied in this regard. I did not understand Stability's concession of "insignificant" incidence to apply in relation to iStock watermarks*.

Models SD XL and v1.6

242. I do not consider there to be any evidence that one real life user in the UK has generated a watermark using either SD XL or v1.6. The Steampunk Musician Image on which Getty Images relies in Annex 8I for SD XL is plainly an image produced during the development of that version of the Model and cannot be attributed to the model weights for SD XL after release. There is no image in Annex 8I for v1.6.

243. The Getty Watermark Experiments failed to produce a result for either of these Models (even using verbatim and re-worded prompts) and the internal Stability Chats show only that the generation of watermarks* was an issue in the pre-release versions of SD XL. That no examples of watermarks* have been found "in the wild" using these Models appears to me to be strongly indicative of the fact that (as the internal Stability Chats suggest) this issue was resolved using appropriate filters applied to the training data and did not recur in respect of either SD XL or v1.6.

244. For reasons I have given above, I do not consider that the generation of one watermarked* image from SD XL (the Donald Glover Image) and three from v1.6 (the Gabba Images) in the Output Claim experiments (all of Getty Images watermarks*) is probative. Ignoring the image which included the word "Getty Images" in its prompt and so is excluded from the scope of the proceedings, two of the images (the Donald Glover Image and the First Gabba Image) were produced using verbatim prompts and one was produced using a re-worded prompt (the second image of The Gabba). As I have explained, there is no evidence whatsoever in real life of such use and thus no basis on which I could find that on balance a real user will have used such a prompt in conjunction with SD XL or v1.6, or that, having done so, a watermark* will have been generated.

245. Ultimately, in respect of these Models, Getty Images invite the court to infer that, given the position in respect of the earlier Models, these later Models must equally have been prone to the production of watermarks*. But to my mind that is not an inference that I may safely draw where (i) there is contemporaneous evidence of steps being taken to address the watermark* issue prior to the release of these models; (ii) there is no evidence whatever of the generation of watermarked* images "in the wild" by these Models after their release – in this context it seems to me to be particularly striking that there is nothing on Reddit given that there is evidence of users of Reddit using the platform to post about watermark* generation in respect of earlier versions of the Model; and (iii) where Professor Brox's expert evidence is to the effect that, given the results of the experiments in relation to the later Models, he can only assume that the filtering has improved.

246. On balance, the fact that it is also the Experts' view that any filter would not be perfect and would leave some images with watermarks in the training data (as is borne out by

the watermarks* on the Donald Glover Image and the Gabba Images) does not affect my view. Absent a clear understanding of how many such images may remain in the training data and what (if anything) the consequence of their presence may be (including whether this remaining number makes it statistically likely that the Model will be prone to overfitting or memorisation such that it will generate watermarks* on synthesised images when subjected to real world usage, or not) I cannot properly find on balance that this has in fact occurred. For present purposes I can only infer on the evidence that the filter used on the training data for SD XL (and I infer also for v1.6) must have been sufficiently effective to preclude the possibility of watermarks* appearing in the wild, other than through the use of contrived and unrepresentative prompts.

247. Getty Images sought to persuade me that the absence of real world evidence of watermarks* being generated is unsurprising and they contend that the situation has analogous difficulties to those which arise in identifying real world evidence of confusion (see by way of example *Fine & Country Ltd v Okotoks Ltd* [2012] EWHC 2230 (Ch) at [84]-[88] per Hildyard J): users of Stable Diffusion are unlikely to report examples of images bearing watermarks* to Stability or to Getty Images because “[t]hey will just assume this is normal behaviour of the model and/or that the model has been trained on photographs from Getty Images”. Furthermore, say Getty Images, images generated using these Models are inherently ephemeral, they are easily producible and easily expendable – thus a user who generates an unwanted image bearing a watermark* is likely to discard it without seeking to draw the issue to anyone’s attention or complain about it.
248. While in general terms I agree with Getty Images on the largely ephemeral nature of the images produced by the Models together with the fact that a user generating an unwanted watermark* is unlikely to do anything other than discard the offending image, nevertheless I do not consider these points to assist on the threshold question. Unlike the issue of confusion, where the likelihood of deception is ultimately a question for the Judge, who is entitled to give effect to his or her own opinion (see *Fine & Country* at [88]), the threshold question in this case requires Getty Images to establish on balance that at least one user in the UK has generated a watermark*. In relation to SD XL and v1.6 (having dismissed the majority of the evidence in Annex 8H as being contrived and unrepresentative), there is not one jot of evidence on which I could properly find on balance that watermarks* have ever been generated in the UK by real world users.
249. The first of the Models in issue has been available since August 2022. Since then, the evidence shows that, by October 2022, DreamStudio had already amassed more than 1.5 million users who, together (according to Getty Images) had created 170 million images. The evidence in Getty Images’ CEA Notice estimates that by August 2023 across all Stable Diffusion variants (including those outside the scope of this claim) the total number of generations had risen as high as 12.59 billion images. Against that background, I agree with Stability that there is a very striking absence of evidence to support Getty Images’ case in relation to the later Models.
250. Ms Cameron’s evidence in cross examination was that Getty Images staff and lawyers had conducted a serious search of online materials for evidence of the generation of watermarks*. They had searched Discord, Twitter and Reddit. Mr Stanley explained that Getty Images’ has a global customer base of 700,000 customers, thousands of

whom would be “enterprise level” and thus in regular contact. 145 million unique visitors to the Getty Images Websites occur each quarter.

251. In my judgment, if evidence of the generation of watermarks* was available, Getty Images had every opportunity to find it. Yet, notwithstanding the “hundreds of thousands, if not millions” of interactions each year described by Mr Stanley as taking place between the global marketing team and customers or potential customers of Getty Images, they have produced no evidence whatever of any real world generation of watermarks using SD XL or v1.6.
252. Getty Images point to correspondence dating back to June 2024 in support of an argument that had a positive case been advanced by Stability on the “threshold question”, disclosure would have been addressed entirely differently. An issue for disclosure was “[i]n what circumstances does Stable Diffusion generate synthetic image outputs bearing watermarks”. In a letter dated 5 June 2024, Bird & Bird recorded that Stability had over 243 million synthetic images (generated by users of DreamStudio) but that these would take 472 hours to transfer out of CoreWeave and that storage of these images would be “incredibly expensive”. Bird & Bird observed that Getty Images had not yet proposed how it would go about analysing this data and said that a disclosure exercise of this magnitude would be disproportionate. Fieldfisher responded on behalf of Getty Images on 25 June 2024 acknowledging that in light of Bird & Bird’s comments, Getty Images would not be seeking disclosure of the Synthetic Images dataset in relation to this (and other) issues for disclosure.
253. Given that it has always been Stability’s case that the generation of images in Annex 8H was contrived, it is unclear why Getty Images decided to take this approach and I reject the suggestion that it would have done anything differently had a positive case been advanced by Stability (as opposed to merely a non-admission). Whether a defendant makes a non-admission or pleads a positive case denying a claim, the burden lies with the claimant to establish its case on the balance of probabilities and Getty Images plainly had the opportunity to analyse (or even to sample) vast quantities of real life data. Whatever Stability’s lawyers may have said about proportionality, Getty Images has always sought to highlight the significance of this case, its importance to Getty Images’ commercial activities and the public interest in its outcome. Getty Images chose not to utilise an available repository of evidence apparently owing to cost and cannot now complain that the outcome might have been different had they taken a different approach.
254. Accordingly I find that the threshold issue is not satisfied in relation to SD XL and v1.6 in respect of Getty Images watermarks*.
255. Furthermore, as was the case with v2.x, there is no evidence whatever that iStock watermarks* have ever been generated in real life using SD XL or v1.6 and accordingly I also find that the threshold issue is not satisfied in this regard.
256. Finally, I note that in opening submissions, both parties made submissions about the issue of “consent”, Stability’s case in its pleadings being that Getty Images have consented to all outputs generated by their agents for the purposes of the proceedings such that those outputs (i.e. in Annex 8H) cannot comprise infringements. During closing, however, it appeared that there was really nothing between the parties on this point. Stability accepts that Annex 8H is relevant to the “threshold” question and I have

found it to be probative in certain respects in connection with that question. I can see no need to go on to consider any further issues on consent and I did not understand the parties to invite me to do so.

257. In light of my factual findings on the threshold issue, the question of trade mark infringement arises only in respect of v1.x (insofar as it was accessed via DreamStudio (for v1.4 only) and/or the Developer Platform) and v2.x. Accordingly, I shall address the question of trademark infringement in relation to these versions of the Model only.
258. Before doing so I observe that, as will already be apparent, the question of trade mark infringement in this case is complicated by the fact that, owing to the stochastic nature of the Models, no two images will be exactly the same and when watermarks* do appear, they will appear in different forms: broadly they may be extremely distorted and/or blurred bearing little real resemblance to the Marks, they may bear only a very loose resemblance to the Marks, or they may look similar, or very similar, to the Marks. This is, of course, very different from the conventional trade mark case in which instances of the offending sign usually look the same, even if the context in which they appear may be different. It also means that the mere fact of an apparent watermark* of some description appearing on a synthetic image may not in itself be sufficient to give rise to a claim, whether under s.10(1), 10(2) or 10(3) TMA. A mere “splodge” or heavily distorted word that might or might not represent an attempt by the Model to generate a watermark* is unlikely to do so, whereas the clearer and less distorted the watermark* becomes the greater the potential for a claim.
259. By way of example, as Getty Images acknowledge, once the watermark* becomes too distorted or blurred, it will not support a claim of trade mark infringement under s. 10(1) TMA. I set out below an image from Annex 8I generated by version 2.1 (“**the Alien Landscape Image**”) in respect of which Getty Images concede that there could be no claim under s.10(1):



The Average Consumer

260. It is well established that many questions in trade mark law are to be assessed from the perspective of the ‘average consumer’ of the relevant goods and/or services, who is deemed to be reasonably well informed and reasonably circumspect (*Lidl Great Britain Ltd v Tesco Stores Ltd* [2024] EWCA Civ 262 (“**Lidl**”) per Arnold LJ at [15])). The average consumer includes “any class of consumer to whom the guarantee of origin is

directed and who would be likely to rely on it, for example in making a decision to buy or use the goods" (*London Taxi Corp v Frazer-Nash Research Ltd* [2017] EWCA Civ 1729; [2018] FSR 7 per Floyd LJ at [34], cited with approval by the Supreme Court in *Iconix Luxembourg Holdings SARL v Dream Pairs Europe Inc* [2025] UKSC 25 ("Iconix") at [29]).

261. In *Lidl* Arnold LJ highlighted various key points as to the identity of the average consumer at [16]-[20] (expressly approved in *Iconix*):

"16. First, the average consumer is both a legal construct and a normative benchmark. They are a legal construct in that consumers who are ill-informed or careless and consumers with specialised knowledge or who are excessively careful are excluded from consideration. They are a normative benchmark in that they provide a standard which enables the courts to strike a balance between the various competing interests involved, including the interests of trade mark owners, their competitors and consumers.

17. Secondly, the average consumer is neither a single hypothetical person nor some form of mathematical average, nor does assessment from the perspective of the average consumer involve a statistical test. They represent consumers who have a spectrum of attributes such as age, gender, ethnicity and social group. For this reason the European case law frequently refers to "the relevant public" and "average consumers" rather than, or interchangeably with, "the average consumer": see, for example, Case C-252/07 *Intel Corporation Inc v CPM United Kingdom Ltd* [2008] ECR 1-8823 at [34]. It follows that assessment from the perspective of the average consumer does not involve the imposition of a single meaning rule akin to that applied in defamation law (but not malicious falsehood). Thus, when considering the issue of likelihood of confusion, a conclusion of infringement is not precluded by a finding that many consumers of whom the average consumer is representative would not be confused. To the contrary, if, having regard to the perceptions and expectations of the average consumer, the court considers that a significant proportion of the relevant public is likely to be confused, then a finding of infringement may properly be made.

18. Thirdly, assessment from the perspective of the average consumer is designed to facilitate adjudication of trade mark disputes by providing an objective criterion, by promoting consistency of assessment and by enabling courts and tribunals to determine such issues so far as possible without the need for evidence...

19. Fourthly, the average consumer's level of attention varies according to the category of goods or services in question.

20. Fifthly, the average consumer rarely has the opportunity to make direct comparisons between trade marks (or between trade marks and signs) and must instead rely upon the imperfect picture of the trade mark they have kept in their mind”.

262. The legal construct of the “average consumer” has been created “to strike the right balance between various competing interests including, on the one hand, the need to protect consumers and, on the other hand, the promotion of free trade in an openly competitive market” (see *Interflora Inc v Marks and Spencer Plc* [2014] EWCA Civ 1403, [2015] FSR 10, (“**Interflora**”) per Kitchin LJ at [113]). In a case concerning ordinary goods and services, the court may be able to put itself in the position of the average consumer without requiring evidence from consumers, still less expert evidence or a consumer survey. In such a case, the judge can make up her own mind about the particular issue she has to decide in the absence of evidence and using her common sense and experience of the world (*Interflora* at [115]).

263. As Arnold J (as he then was) observed in *Sky Plc v SkyKick* [2018] EWHC 155 (Ch) at [275], “[t]he average consumer for the purposes of an infringement claim must be a consumer of the relevant goods and/or services who is both (i) familiar with the trade mark and (ii) exposed to and likely to rely upon, the sign”.

264. In practice in the present case, Stability admits that Getty Images has a substantial reputation and goodwill as a creator of digital visual content in the UK and as a provider and licensor of such content to customers in the UK. Accordingly, it can safely be assumed that all the potentially relevant consumers are familiar with it. However, the class of average consumer who is “exposed to, and likely to rely upon, the sign” must be considered and, as both parties accept, that class may itself be split into different classes of average consumer with different attributes depending upon the different ways in which each class will encounter the signs in issue on a synthetic image output from Stable Diffusion.

265. By closing submissions, it was common ground therefore that there are at least three categories of average consumer in this case namely:

i) the average consumer who downloads the codebase and model weights from GitHub and Hugging Face respectively and runs the inference offline (whether using the Direct Download method or the Diffusers Method). As the evidence of Professor Farid, Ms Gagliano and Ms Kabba-Kamara confirms (and as I find), this average consumer is technically skilled and has access to a high powered computer (albeit, as Professor Farid confirmed, in his experience, any modern laptop or desktop would fall into this category). Downloading the Model enables the user to make bespoke changes to it. Ms Gagliano’s evidence explains that the Models are often used to power text to image functionality in various products (such as design tools, stock sites such as 123RF, or lesser known startups), and that it is “very common” for these companies to build additional applications on top of Stable Diffusion models. Professor Brox explains that users can ‘branch’ from the provided code by using third party code (which amends Hugging Face) or by writing their own code. The Model Cards for all versions of Stable Diffusion evidence the technical sophistication that a consumer adopting this method of access needs to have.

ii) the average consumer who runs the inference on Stability's computing infrastructure using an API, or in other words, (as it now is) the Developer Platform, found online at platform.stability.ai. As the evidence of Professor Farid and Ms Kabba-Kamara confirms (and as I find), this average consumer has a good level of technical knowledge, albeit less than the consumer who downloads the Model. This consumer must access the Stability website and create an API account with Stability. This enables him to use appropriate code running on his computer (or on the AWS Instance, as used by Getty Images in conducting its experiments) to send queries in a compatible format to Stability over the internet – Stability then runs the inference using its own computing infrastructure and returns the resulting images to the user.

iii) the average consumer who uses the Stability DreamStudio service to run the inference on Stability's computing infrastructure using a web interface, which Ms Gagliano explains is like a webpage. This requires the user to access the DreamStudio website at beta.dreamstudio.ai through a normal browser, such as Chrome, Safari or Firefox, and to create an account having agreed to the Terms of Service. The account creation service involves an email sign-up and confirmation of a verification email sent to the user's email account. Ms Gagliano accessed DreamStudio by logging in with her Discord account and then paying for access by purchasing credits. The user must then select the Model (now only SD XL and v 1.6 are available on DreamStudio) and the process of inference is then commenced by clicking a button marked 'Dream'. The user is given the option of how many images to generate at once, the default number being four. I accept Ms Gagliano's evidence (which is confirmed by the exhibits to Ms Kabba-Kamara's statement) that using DreamStudio is a more "user friendly" way to access the Models and so is available to consumers with access to a computer and a degree of familiarity with using the internet, together with a willingness to agree to the Terms of Service. Professor Farid expressed the view, which I accept, that this user requires "no special technical knowledge".

266. In addition, however, Getty Images now contend that ordinary members of the public (i.e. consumers with no, or very little in the way of, technical skills) will be exposed to the Sign on synthetic images generated by Stable Diffusion because (i) they have asked someone else to download the Model for them before using it on their own computer to generate an image; or (ii) they have been shown an output image from the Model by someone else, for example a friend or work colleague (Getty Images describe this as "post-sale context"). Getty Images therefore submits that this is another, distinct class, of average consumer which the court should consider.

267. I must take care with this contention. Not only have these different ways of being exposed to the Sign not been pleaded (as strictly they should have been so as to alert Stability to the case it needed to meet at trial), but various of the means by which it is suggested by Getty Images that the average consumer will encounter the sign also appear to me to be speculative.

268. As Mr Cuddigan submitted in closing submissions, I consider it to be unrealistic to posit a scenario where a technically savvy third party will download the Model onto the computer of a technical ingénue, because even operating Stable Diffusion once it has been downloaded requires considerable technical ability together with a computer

powerful enough to do so. I can see no reason why someone with the technical ability to operate Stable Diffusion once it is downloaded would need anyone else to download the Models for them, and none was suggested. I have seen no evidence to suggest that this is something that occurs in the real world. As Mr Cuddigan said, “[y]ou have DreamStudio for people like that”.

269. In closing, Getty Images appeared to suggest the possibility (based on evidence from Mr Stanley) that designers might “screen grab” images with watermarks* on them in order to mock up a design which will be seen by others. Mr Stanley explained that the presence of the watermarks* will enable the designers to identify the source of the image and obtain a licence for it later, if required. However, while I accept Mr Stanley’s evidence about this, it does not assist when one is talking about the use of synthesised images which require no licence. It is not realistic to suppose that a designer using Stable Diffusion to synthesise an image would need to take a screen grab of a watermarked* image in order to use it as an aide-mémoire as to its source. Such a designer knows he has obtained the image using Stable Diffusion, that he does not need a licence for it and, if it includes a watermark*, that watermark* will, in any event, be unwelcome. He will certainly not be using it to remind himself, or others, of the source of the image.
270. Perhaps more difficult, however, is the question of whether it is realistic to suppose that users of Stable Diffusion will forward images bearing watermarks* to others, or (possibly) share their screen showing such images with others in the ‘post-sale’ context.
271. I pause to acknowledge that there is inevitably an overlap between this issue and the issue of context to which I shall come in a moment. However, I take the issue of ‘post-sale context’ at this stage so as to address Getty Images’ submissions on the existence of an additional category of technically unsophisticated average consumer.
272. There is no dispute on the legal principles to be applied. As recently confirmed in *Iconix* at [40], when the court is considering post-sale context it is concerned with the “realistic and representative way in which the average consumer will encounter the sign”.
273. While there may be many trade mark cases in which the court’s assessment of post-sale context is uncontroversial and needs no supporting evidence, perhaps because the sign was invisible at the point of sale (as with the mark traditionally placed upon the cork in a bottle of wine) or because (although the average consumer did not purchase the product) he or she admires it being worn by a friend or acquaintance (see *Montres Breguet* at [84] and *Iconix* at [40]), the facts of this case are not so straightforward. As Getty Images’ pleaded case appears to acknowledge, the realistic and representative way in which the average consumer will encounter the Sign is likely to be immediately following the generation of an image which then appears on his or her computer screen.
274. I start from the premise that it must certainly be the case that users of Stable Diffusion will, on (many) occasions, share the images they generate with others – at least in the work context that may very well be the purpose of generating the images. Equally, however, as is common ground, those users will not be trying to generate images which bear watermarks*. Stable Diffusion is used by people who want to generate images for use in a presentation or a project or an artwork, as Ms Varty confirmed. They want images that “look good” and, even assuming that Stable Diffusion is also used for

leisure purposes (i.e. “for fun”), it is likely that such users will have the same aim in mind – at least they will want images that correspond to their text prompts. They will not want their generated images to be marred by the presence of a watermark*. I consider it to be reasonable to infer that it is images that “look good” and respond most accurately to the imagination of the user as reflected in his or her text prompt, that the user will normally wish to generate and (potentially) to share.

275. The Reddit Exchanges illustrate that users will be irritated if they generate an image that is spoilt by the presence of a watermark of any kind and will want to get rid of it. The user with the tag namesareunavailable describes the experience of generating watermarks as “quite annoying”, while Antique_Plane_130 asks: “What’s the best way to remove a watermark from an image”. Suggestions in the various exchanges include changing words, or using specific phrases, in the prompt or trying to use “negative weight prompts”; or using a model “trained to remove watermarks”. There is no evidence that anyone wants to generate images with watermarks*. As I have already indicated, Getty Images posits a case in its pleadings that the average consumer may use Stable Diffusion “to generate an image which is the same as or similar to the Claimants’ Content but in respect of which it will not have to pay a licence fee”. But, even assuming this to be correct (and as I have said there is no evidence to support this proposition), it makes no sense that such a user would be trying to produce a watermarked* image or that he or she would then want to make any use of such an image, or show it to someone else.
276. The only evidence of users sharing watermarked* images appears to be with a view to pointing out the anomaly of the watermark* appearing. Again this is clear from the Reddit Exchanges. These show a handful of users sharing watermarked* images and/or their experience of generating watermarked* images with other (often well-informed) users on a forum devoted to discussion about the Models. This appears to be in circumstances where there is a common understanding across the vast majority of these users that this is unusual or extraordinary and that the watermarks* are a consequence of the data used to train the Model.
277. Given the evidence, I do not consider that I can find (as Getty Images invite me to do) that it is realistic and representative to suppose that there is a class of unsophisticated average consumer who will be shown images generated by Stable Diffusion which bear a watermark* by someone else. Indeed, Getty Images’ submissions to this effect appear to me to be somewhat inconsistent with their submissions (recorded earlier in this judgment) that watermarked* images are likely to be discarded as unwanted by users.
278. Further, if Getty Images wished to advance such a case then it should have been pleaded and the circumstances in which they said it was likely to occur should have at least been identified. I note in this context that this particular scenario was mentioned purely in one line in Getty Images’ opening skeleton argument for trial. There was nothing in the List of Issues to even hint at the possibility that this was a serious point that Getty Images was intending to advance.
279. Returning to the three categories of average consumer identified above, I accept Stability’s submission that it is reasonable to infer that the first two categories of average consumer have a familiarity with AI models. I consider the evidence to show, and I find, that a significant proportion of these average consumers will understand how the Model was trained and that the generation of watermarks* is a function of the data

on which the Model was trained. Users in the first category will also know that the Model they are running is operating on their local GPU. However, I reject Stability's case that these two classes of consumer will understand that Stability has "no control" whatever over the images that are generated, or to put it another way, that they (the users) have complete control over those images, whatever is generated – a point which appears to me to be inconsistent with the evidence in the Reddit Exchanges to which I have referred above.

280. The average consumer using the Model via DreamStudio will not necessarily, in my judgment, have a similar level of understanding of AI models. A significant proportion will not appreciate how the Model has been trained or why it might produce watermarks*. I did not understand Stability to suggest otherwise.
281. Getty Images accept that the degree of care and attention that may be exercised by these various categories of average consumer may be different and contends that in relation to "business consumers" (by which I understood them to mean consumers accessing the Model for a particular purpose connected with business) that level of attention may be moderate, whereas in relation to "members of the general public" (by which I understood them to mean anyone else who was likely to use Stable Diffusion to generate an image) will pay only a low degree of attention. Getty Images pray in aid Mr Stanley's evidence to the effect that Getty Images' customers "who are busy people" may not notice details in a watermark itself.
282. Getty Images did not seek to identify how these submissions are to be applied to the three categories of average consumers, but I can only assume that their case is that within each category there will be both business users and "others", whose levels of attention will be different.
283. Stability rejects this characterisation of the degree of attention that an average consumer would pay, arguing that the average consumer using any one of the three access mechanisms (whether for business or otherwise) would pay "the highest degree of attention", because the very purpose of using the Model is to generate bespoke images in response to a user created text prompt.
284. In my judgment, while some average consumers (who are seeking to generate an image for use in a work context such as a presentation) may indeed pay a high degree of attention to the synthetic images generated (as Mr Cuddigan elegantly put it "[t]hey are curating through their prompts an extraordinarily powerful image creation technology"), I consider that not all would do so. This applies to all three access mechanisms and is (of course) dependent upon the purpose for which the images are being generated.
285. I reject Getty Images' submission that any class of average consumer will ever give only a low degree of attention to images generated by the Model. Anyone using the Model to generate bespoke images in response to text prompts is likely, in my judgment, to give those images at least a moderate degree of attention, whatever the intended use of the images and notwithstanding that he or she has had to pay nothing, or only very little, for them. I do not consider that this conclusion is affected by Mr Stanley's evidence. Getty Images' customers do not merely comprise individuals using their GAI to generate images and I do not consider Mr Stanley's evidence to assist on the level of

attention that a user of Stable Diffusion would give to the images that he or she generates when using the Model.

Context

286. It is common ground that the court must consider the actual use of the sign complained of in the context in which it has been used. As Kitchin LJ (as he then was) said in *Specsavers International Healthcare Ltd v Asda Stores Ltd* [2012] EWCA Civ 24 (“Specsavers”) per Kitchin LJ at [87]:

“In my judgment the general position is now clear. In assessing the likelihood of confusion arising from the use of a sign the court must consider the matter from the perspective of the average consumer of the goods or services in question and must take into account all the circumstances of that use that are likely to operate in that average consumer’s mind in considering the sign and the impression it is likely to make on him. The sign is not to be considered stripped of its context”.

287. However, the parties are not in agreement over the application of this principle. Getty Images submit that the court must avoid an over-expansive view of context which might turn an otherwise infringing use into a non-infringing use. They contend that context in the circumstances of this case is limited to what is on the screen when the image bearing the watermark* appears and they rely in particular upon the observations of Arnold LJ in *Lifestyle Equities v Royal County of Berkshire* [2024] EWCA Civ 814 [2024] FSR 32 (“**Lifestyle Equities**”) at [51]-[58]. Arnold LJ begins at [51] by explaining that the use of the allegedly infringing sign must be considered in context. He then goes on to say:

“52 This principle is perhaps most clearly illustrated by Case C-533/06 *O2 Holdings Ltd v Hutchison 3G UK Ltd* (C-533/06) [2008] E.C.R. 1-4231. O2 was the proprietor of trade marks consisting of images of bubbles registered in respect of telecommunication services. H3G used similar images of bubbles in respect of identical services in a comparative advertisement. The CJEU held that the bubble images had been used by H3G directly in relation to O2’s services and indirectly in relation to H3G’s own services. That gave rise to the question of whether the latter use was infringing use. The Court held that there was no likelihood of confusion because the national court had found that the comparative advertisement as a whole was not misleading and did not suggest that there was any form of commercial link between O2’s services and H3G’s services. The Court stated at [64] that “the referring court was right to limit its analysis to the context in which the sign similar to the bubbles trade marks was used by H3G, for the purpose of assessing the existence of a likelihood of confusion”. It also stated at [67] that “the assessment must be limited to the circumstances characterising that use”. This suggests that the assessment is limited to what might be termed the immediate context in which the sign has been used. The Court nevertheless implicitly

rejected the argument that distinguishing context was irrelevant if the signs themselves were confusingly similar to the trade marks and used in relation to identical services.

53 The decision in *O2 Holdings Ltd v Hutchison 3G UK Ltd* (C-533/06) was considered in some detail by Kitchin LJ in *Specsavers* at [80]-[84], and formed the basis for his statement of principle at [87] which I have set out in paragraph 12 above.

54 Although the principle is clear, its application can cause difficulty. The difficulty usually arises where the defendant relies upon context as negating a likelihood of confusion in a case where, absent whatever is relied upon as constituting the relevant context, the identity or similarity of the mark and the sign and the similarity or identity of the respective goods or services would give rise to a likelihood of confusion. While it is clear from *O2 Holdings Ltd v Hutchison 3G UK Ltd* (C-533/06) that this is legally possible in an appropriate case, it is not clear how far the principle extends outside the special circumstances of comparative advertising.

55 The judge cited two passages from the judgment of Daniel Alexander QC sitting as a Deputy High Court Judge in *PlanetArt LLC v Photobox Ltd* [2020] EWHC 713 (Ch); [2020] F.S.R. 26 which addressed this question. After citing *Och-Ziff Management Europe Ltd v Och Capital LLP* [2010] EWHC 2599 (Ch); [2011] F.S.R. 11 and *Specsavers*, Mr Alexander QC went on:

"24. In my view, Arnold J, as he then was, in *Och-Ziff Management Europe Ltd v Och Capital LLP* was saying that the CJEU took the view that, in considering infringement of a registered trade mark, it was not appropriate to look so broadly at the context that use which was *prima facie* infringing was nonetheless to be regarded as non-infringing because other, separate, acts of the defendant had countered actual deception. An extreme example is where a defendant uses a well-known brand for counterfeit goods but nonetheless makes it very clear that the goods are in fact counterfeit so that no actual purchaser is confused. There may be no actual confusion as a result of the use of the sign but there is nonetheless trade mark infringement because the court must focus on the use of the sign in question not the other statements by the defendant as to the trade origin of the goods.

25. Accordingly, while it is right to take the context in which the given sign will be seen into account, I am not persuaded that it would be right to expand the view so broadly as to take account of the fact that a given sign only appears in this case after a different sign has been used. To that extent, each use

of the signs must be examined separately in what might be described as its 'local' context.

...

164. I have also considered, in accordance with the guidance in *Specsavers International Healthcare Ltd v Asda Stores Ltd*, and generally whether there is anything about the context of presentation of the marks which negates that result. In my view there is not. It is also necessary for the court to be cautious in adopting an overly expansive approach to taking account of context in a trade mark claim. One purpose of registered trade mark protection (in which it is distinguished from passing off) is to provide an element of exclusivity in the use of a registered mark, regardless of the wider context in which it is used, so long as the conditions for protection are fulfilled."

288. At [56] and [57] Arnold LJ points out that there are other answers to the problem identified by Mr Alexander at [24]. At [58] he concludes:

"58. Having cited *PlanetArt LLC v Photobox Ltd*, the judge said at [67] that "[t]here are sound policy reasons for not taking an over-expansive view of the context of the allegedly infringing use". While I have some sympathy with that point of view, the issue of how far context extends is a difficult one for the reasons explained above. It follows that it is best decided in a case where it actually matters. It does not matter in this case, because the crowded market is relevant to the distinctive character of the Trade Marks regardless of how narrowly or broadly the context of the allegedly infringing use is drawn".

289. Further, Getty Images drew my attention to *Shorts International Ltd v Google* [2024] EWHC 2738 (Ch) [2025] ETMR 3 ("**Shorts International**"), per Michael Tappin KC, sitting as a Deputy High Court Judge. At [186] Mr Tappin KC identified that a question in that case was whether context can include matter which is not visible to the user at the time at which the sign is used. He then set out various observations made by Arnold J in *Och-Ziff Management Europe Ltd v Och Capital LLP* [2010] EWHC 2599 (Ch) ("**Och-Ziff**") (including at [77] to the effect that "the context and circumstances are limited to the actual context and circumstances of the use of the sign itself", that circumstances are "circumstances characterising the use" and not circumstances more generally and that "circumstances prior to, simultaneous with and subsequent to the use of the sign may be relevant to a claim for passing off...but they are not generally relevant to a claim for trademark infringement"), before referring to the passages in *LifeStyle Equities* which I have quoted above. At [189] he said this about those passages:

"It appears to me that Arnold LJ was expressly not approving (or disapproving) what Mr Alexander said in the passages from *PlanetArt* which he cited, and leaving open the question of the proper scope of "context". I propose to proceed on the basis that

the law permits, and indeed requires, the court to consider each use of the sign taking into account the context in which consumers will see the sign, but that the law does not permit the court to take into account matter which is not visible to the consumer when that use of the sign is being made. That is particularly the case where not all consumers will have been exposed to such matter”.

290. Mr Tappin KC applies this conclusion to the facts of that case at [227] where he takes into account as “part of the context of the use” the fact that the YouTube name and logo is always present on the screen when a particular sign appeared; and at [229], deciding that, by contrast, he should not take into account, as part of the context of use of the sign “the fact that the user accessed the Shorts feed through the YouTube app and via the home page of the app. In any event...it is entirely realistic that some consumers will be shown another’s screen with the Shorts feed already open”.
291. Stability submits that the passages to which I have referred in the authorities say no more than that the question of context must be considered on the facts of the case and is an open question. I agree. The touchstone, as is clear from *Specsavers*, is to identify the circumstances that are likely to operate in the mind of the average consumer when considering the sign and the impression it is likely to make on him or her. Nothing said by the judge in *Shorts International* appears to me to contradict this proposition. On the facts of that case, the court found that there was a realistic category of consumers who would be shown the sign without having gone through the prior stages within the app, which had been alleged to provide relevant context. I have already determined that the only three categories of average consumer in this case (as pleaded) would all have accessed the Model themselves via one of the three access mechanisms and accordingly the key question for me will be the relevant context in respect of each of those access mechanisms.
292. In this regard, I again agree with Stability that the Court of Appeal in *Specsavers* did not place any temporal limit on relevant context but, rather, specified “all the circumstances” likely to operate in the average consumer’s mind. I also note that in a passage at [86] of his judgment in that case, Kitchin LJ observed that it was not clear to him what Arnold J and the parties had in mind in *Och-Ziff* by the phrase “circumstances prior to, simultaneous with and subsequent to the use of the sign” and he went on to explain that the phrase was to be seen “in light of the particular and rather specific issue in that case”. It was against this background that Kitchin LJ then set out the “general position” in [87] to which I have referred above. The limitation of the phrase used by Arnold J in *Och-Ziff* to the facts of that case is entirely consistent with the Court of Appeal recognising that “all the circumstances” likely to operate on the average consumer’s mind must be considered. Indeed, at first instance in *Specsavers* ([2010] EWHC 2035 (Ch); [2011] FSR 1), Mann J held that a circumspect consumer encountering the Specsavers sign in an Asda store will know that he is in that store, just as the consumer who encounters the Specsavers logo online will know that he or she has entered an Asda site already (*Specsavers* at [137]). These findings were not challenged in the Court of Appeal. I note also Arnold LJ’s observation in *Montres Breguet SA v Samsung Electronics Co Ltd* [2023] EWCA Civ 1478, [2024] ETMR 13 (“*Montres Breguet*”), at [76] to the effect that it is not necessary to confine attention to

the “immediate” context in which the signs were used. Everything will depend on the facts of the particular case.

293. Given their differing approaches to the question of context, the parties do not agree on whether the various features of each of the three access mechanisms in this case are relevant to the court’s consideration of context, including whether the average consumer would read information/terms and conditions set out in the Model Cards, on the GitHub and Hugging Face pages, the Stability AI API Terms of Service and in the Terms of Service for a DreamStudio account. I set out in Appendix B to this judgment the (largely uncontroversial) factual circumstances encountered by the average consumer in accessing version 1.x (via DreamStudio and the API) and v2.x of the Model (via each of the three access mechanisms).
294. Getty Images contends that Stability’s analysis of these factual circumstances involves considering “an artificial context in which the average consumer would interact with the sign for the purposes of negating infringement”, that it is wrong in law and that, in any event, it falls down on the facts.
295. My conclusions as to the relevant context in which watermarks* would be encountered by the average consumer using each of the three access mechanisms are set out below.

Downloading from GitHub/Hugging Face (v2.x only)

296. Stability submits that the features set out at Appendix B, including the Model Cards and the License terms, form part of the relevant context in the *Specsaver* sense, because a reasonably circumspect average consumer would have regard to them prior to downloading the Model.
297. I accept that the technically sophisticated average consumer accessing the Model by this means is likely to read the Model Card together with the GitHub and Hugging Face pages (not least because these set out useful and important information about how to get the most out of using the Model and will also be relevant to the user’s choice of Model). While there is no reason to suppose that this will happen at any time *after* the Model has been downloaded onto a local computer (at which point the Model Card is no longer visible), I do not consider that one can properly ignore the general understanding that the average consumer will have gleaned from this information and it would be artificial to do so.
298. I consider that the technically sophisticated average consumer using this means of access will at all times understand from his interaction with the Model Card and the GitHub and Hugging Face pages that, once he has downloaded it, he is using a cutting edge AI model to generate synthetic images using his own prompts and that the Model is not designed to produce out-of-scope images i.e. factual or true representations of people or events (whether he chooses to try to do so or not). He will also understand that the image generated on his own computer has been generated without the intervention of Stability and that downloading the Model enables him to add bespoke elements to it. I have already said that this average consumer (owing to his level of technical sophistication) will understand how the Model is trained (and this is also addressed in the Model Cards) and why watermarks* may be generated.

299. Stability also contends that the terms and conditions of the License are relevant context, because the average consumer using this access mechanism would read those terms and conditions. Specifically Stability submits that, given the rationale for the average consumer's characteristics of striking the right balance between the protection of consumers and the promotion of trade, there can be no justification in protecting the incautious – hence the requirement that the average consumer is “reasonably circumspect”. Apparently in recognition of the fact that (as Getty Images contend) many people will not routinely read terms and conditions, Stability says that the “reasonably circumspect” test must be applied “regardless of how many consumers might achieve that standard in the real world” and it prays in aid *Morley's (Fast Foods) Ltd v Nanthakumar* [2025] EWCA Civ 186 (“*Morley*”).

300. In *Morley*, the Judge at first instance had identified a class of average consumer for a chicken shop as including late night and early morning revellers, a significant subset of whom she said “will be intoxicated”. On appeal, Arnold LJ held that the Judge had been wrong in her identification of the class of average consumer and that she had wrongly focused on consumers who were intoxicated “and therefore by definition not reasonably well informed and circumspect” (at [16]-[17]). Stability says that, by analogy, this court should not have regard to the fact that many consumers may not read terms and conditions. It points out (without reference to any particular dictionary definition) that the ordinary meaning of the word ‘circumspect’ is someone who is “careful not to take risks” or someone who is “showing caution, cautious, wary: taking everything into account” and it says that such a consumer would not proceed to generate images using Stable Diffusion “without regard to the terms on which they did so”.

301. On balance, in the case of this particular access mechanism, I disagree with Stability that the average consumer would read the License terms when accessing the Model.

302. I did not understand Getty Images to dispute the meaning of the word ‘circumspect’ and I accept that it is intended to encompass a careful consumer who does not take risks. The question here, however, is where one sets the bar. On the facts of the present case, that depends on whether the reasonably well-informed and circumspect average consumer will read terms and conditions in circumstances where (i) they are not set out on the web pages from which the source code and weights for the Model can be downloaded; (ii) it is not anywhere suggested on those web pages or in the Model Cards that users wishing to download the Model must read the terms and conditions set out in the License; (iii) it is necessary to click on a hyperlink to see the License; (iv) even if a user does use the hyperlink to reach the License, the terms and conditions are extremely dense; and (v) there is no specific requirement (or even recommendation) anywhere on the web pages or in the Model Cards for a consumer to confirm that the terms and conditions of the License have been read and/or accepted prior to downloading the Model.

303. In my judgment, while some particularly rigorous consumers will no doubt read the License, a significant proportion will not. I do not consider that, in the circumstances of this case, this significant proportion will fall outside the category of the reasonably well-informed or circumspect consumer and I do not consider there to be any sensible analogy to be drawn with “intoxicated” consumers who could not in any circumstances be described as ‘circumspect’. In her closing reply submissions, Ms Lane referred me to *Schutz (UK) Ltd v Delta Containers Ltd* [2011] EWHC 1712 (Ch), in which Briggs J (as he then was) considered a “reasonable but less than rigorous inspection” of

products would not have led the average consumer to understand the message conveyed by disclaimers which were not prominently printed on labels affixed to the products. While Briggs J did not there consider whether the average consumer would have read the disclaimers, I obtain some support from this decision for the distinction between “a rigorous” and “a reasonable” approach.

304. Having regard to the points I have made, the circumspect consumer (who I have already found will read the Model Cards and web pages in full and so will be reasonably well-informed at the point of downloading the Model) will not consider it a “risk” to proceed without reading the License terms – indeed their familiarity with AI models of this type is likely to confer confidence that absent any specific injunction to read the terms and conditions of the License, there is no need to do so. They will instead conclude from the content of the web pages and relevant Model Card that the key features of use of the Model of which they should be aware are there set out and that it is not a requirement of use that they read or accept the License.
305. Ms Varty did not herself download the Models by this means but she confirmed that the individual who downloaded the Models for her had not drawn her attention to any terms and conditions. There is no evidence from Ms Kabba-Kamara that she (or anyone else) read the License terms before conducting inference.
306. During closing submissions, I understood Mr Cuddigan to say that the real point he was making in relation to this access mechanism was that average consumers who download the Model using this means will understand from the context that they are “using a tool on their account”. If this simply means that they will understand that they are generating images by using their own bespoke prompts on their local machine then I do not disagree and I did not understand Getty Images to disagree either. I consider this to be a function of the sophistication of the average consumer using this access method. However, if this was intended to suggest that such a consumer considers himself or herself to be in total control over, and entirely responsible for, the output of the Model, including the generation of watermarks* (whatever they may have gleaned from their reading of the relevant materials), then I reject that submission for reasons to which I shall return later in this judgment.
307. Finally, in so far as the Diffusers Method requires use of further software in order to set up the Model on a local computer, I agree with Getty Images that these steps have nothing whatever to do with the context in which the user will see the Sign on the synthetic image generated.

DreamStudio

308. Stability accepts that (as I have found) the DreamStudio platform makes Stable Diffusion available to less technically sophisticated users but it points out that this is contingent upon active agreement to the Terms of Service. Further it points out that the DreamStudio user is repeatedly presented with a series of trade marks which indicate that the origin of the service provided is Stability. This occurs at registration, at login, in the Terms of Service, in the web address and at the head of the page when images are generated.
309. I accept that these are locations where the average consumer would expect to find an indication of origin of the online service which he or she is using, as Mr Stanley

confirmed in his evidence and also that western users will start to read a page from the left hand side. The images appear in a section of the page that is dedicated to outputs of the Model, directly under and by reference to the prompt which has been chosen and input by the user.

310. Further, Stability submits that users are repeatedly reminded that they are using a service which generates artificial images:

- i) The service is called DreamStudio;
- ii) The options at the top of the sidebar show that the selected mode is ‘Generate’;
- iii) The user is invited to select model, generation steps, and other parameters; and
- iv) The process of generating images is triggered by the user pushing the button marked ‘Dream’.

311. Mr Stanley accepted that the corresponding branding on the GAI was intended to make clear that the service was producing AI outputs. In my judgment, the average consumer using DreamStudio will at all times understand that he is on Stability’s website, that the images are generated by an AI model and that the Model produces the images in response to his own prompts. Many consumers will also understand from the context to which I have referred that the Model is designed only to generate “artificial images”, although I consider that a significant proportion will not.

312. In my judgment, the average consumer, who is required to accept terms and conditions before accessing the Model, will also have read, or at least skim-read, the Terms of Service before accepting them. Ms Gagliano said under cross examination that she had not read the terms “in detail”, but she did not say that she had not looked at them – and that is what I would expect in circumstances where a consumer is required to accept the Terms of Service before using the Model. Given that the Terms of Service say that the user “must review and follow the terms of the License”, I also consider that the average consumer will have read, or at least skim-read the License. However in my judgment the average consumer is unlikely to recall the detail of the Terms of Service or the License or necessarily to understand their potential significance in connection with the generation of any particular synthetic image. Specifically, I reject Stability’s case that the average consumer using DreamStudio will understand that any attempt to generate a photorealistic image is ‘out of scope’ for the Model. Those who understand that it is designed only to create artificial images may understand this, but a significant proportion will not. In this context I note that the Terms of Service contain no reference to “out of scope” use.

The Developer Platform

313. As I have found, an average consumer using the Model via this means is technically skilled. For similar reasons to those that apply for the DreamStudio access mechanism, he or she will, in my judgment, have read, or at least skim-read, the Stability AI API Terms of Service prior to accepting them, although I do not consider that the average consumer will have read the License terms. Unlike the Terms of Service applicable to the use of DreamStudio, there is no requirement to do so in the API Terms of Service and, accordingly (for similar reasons to those given above in relation to the download

mechanism) the average consumer (who for this access mechanism is technically savvy) who has looked at the API Terms of Service will not consider it a risk not to examine the License terms. Some rigorous consumers may do so but a significant proportion will not. The average consumer who has read the API Terms of Service is unlikely to recall the detail of those terms or appreciate their potential significance in connection with the generation of any particular synthetic image.

314. Although there is no evidence of any specific name or logo being visible when images are generated using this means of access, I consider that the average consumer using this access mechanism will understand that the images are generated by an AI model provided by Stability and that the Model produces the images in response to their own prompts. The average consumer will also understand that the images produced are artificial. I reject the suggestion that the average consumer will understand or remember that attempting to generate photo-realistic images is ‘out of scope’ for the Model, or what, if anything, the significance of that might be. Many may, but a significant proportion will not. The API Terms of Service make no direct reference to out of scope use.

SECTION 10(1) INFRINGEMENT

The Law

315. Section 10(1) TMA provides that:

“A person infringes a registered trade mark if he uses in the course of trade a sign which is identical with the trade mark in relation to goods or services which are identical with those for which it is registered.”

316. Section 10(1) thus provides for infringement absent any likelihood of confusion (which is required under section 10(2)), but it depends on identity between the mark and the sign and identity of goods and services.

317. The law on infringement under section 10(1) was summarised by Arnold LJ in *Easygroup Ltd v Nuclei Ltd* [2023] EWCA Civ 1247 (“*Easygroup*”) at [48] by reference to six conditions, namely:

“(i) there must be use of a sign by a third party within the relevant territory; (ii) the use must be in the course of trade; (iii) it must be without the consent of the proprietor of the trade mark; (iv) it must be of a sign which is identical to the trade mark; (v) it must be in relation to goods or services which are identical to those for which the trade mark is registered; and (vi) it must affect, or be liable to affect, one of the functions of the trade mark”.

318. In this case there are issues as to (ii), (iv), (v) and (vi). The burden of establishing (ii), (iv) and (v) lies with the trade mark proprietor. However, as Ms Bowhill submitted in closing, the claimant in a section 10(1) case benefits from a rebuttable presumption of likelihood of confusion (see *Bentley 1962 Ltd v Bentley Motors Ltd* [2020] EWCA Civ 1726; [2021] FSR 14 per Arnold LJ at [30]-[34] and *Match Group LLC v Muzmatch*

Ltd [2023] EWCA Civ 454; [2023] FSR 18 (“**Muzmatch**”), per Arnold LJ at [116]). Stability has not sought to advance any pleaded case that, notwithstanding any double identity that may be established, there is no effect on the function of the trade mark such that the presumption should be rebutted.

319. I did not understand the applicable legal principles to be controversial, although their application to the facts of this case is in dispute. I shall set out the legal principles in the following section before dealing with their application when conducting the global assessment.

Use of a Sign

320. As to ‘use’, the TMA provides a list of specific activities which are deemed to be use of a sign, at 10(4):

“For the purposes of this section a person uses a sign if, in particular, he—

- (a) affixes it to goods or the packaging thereof;
- (b) offers or exposes goods for sale, puts them on the market or stocks them for those purposes under the sign, or offers or supplies services under the sign;
- (c) imports or exports goods under the sign;
- (ca) uses the sign as a trade or company name or part of a trade or company name;
- (d) uses the sign on business papers and in advertising
- (e) uses the sign in comparative advertising in a manner that is contrary to the Business Protection from Misleading Marketing Regulations 2008.”

321. As Arnold LJ observed in *Montres Breguet* at [56] (with reference to various CJEU decisions, which I have removed):

“The CJEU has repeatedly held that a party only “uses” a sign for this purpose if it uses that sign “in its own commercial communication...A person may allow its clients to use signs without itself using those signs...Merely creating the technical conditions necessary for the use of a sign and being paid for that service is not sufficient to amount to use...”.

322. For a claim to trade mark infringement to succeed, there must be some “active behaviour” on the part of the defendant “and direct or indirect control of the act constituting the use”, such as ‘affixing’ the sign on the goods, or ‘offering’ or ‘supplying’ services under that sign. Acts carried out by an independent operator without the consent of the defendant or against his express will cannot be attributed to a defendant as an infringing use of the sign. (See C-179/15 *Daimler AG v Egyud Garage* EU:C:2016:134 (“**Daimler**”) at [39]-[40] and *Coty Germany GmbH v Amazon Services*

Europe SARL (ECJ) [2020] Bus LR (“**Coty**”) at [37] and [38]). An e-commerce platform like eBay is not responsible for any infringing use where its users list counterfeit or grey market goods for sale on eBay’s website – it merely enables that use to be made by its customers (see C-324/09 *L’Oréal SA v eBay International AG*; [2012] Bus LR 1369 (“**L’Oréal v eBay**”) at [103]-[105]). A service provider who creates the technical conditions necessary for the use of the infringing sign and is paid for that service will not without more have any legal responsibility itself for use of the sign (see *Coty* at [43] and *Google France Sarl v Louis Vuitton Malletier SA* (C-236-238/08) [2011] Bus LR 1 (“**Google France**”) at [52]-[57]).

323. The CJEU has held in a long line of cases that use “in the course of trade” requires (i) use in the context of a commercial activity with a view to economic advantage; and (ii) use that is a non-private matter (see C-206/01 *Arsenal Football Club v Reed* (ECJ) [2003] Ch 454 (“**Arsenal Football Club**”) at [40]). The meaning of this test has been the subject of debate; specifically as to whether the words “not a private matter” impose a separate criterion or are simply intended as providing a contrast to the words “in the context of commercial activity” (see *Och-Ziff Management Europe Ltd v Och Capital LLP* [2010] EWHC 2599 (Ch); [2011] FSR 11, per Arnold J at [56]-[66] and *APT Training & Consultancy Ltd v Birmingham & Solihull Mental Health NHS Foundation Trust* [2019] EWHC 19 (IPEC), [2019] ETMR 22, per HHJ Melissa Clarke at [29]-[39]). It is not clear to me that I need to grapple with this debate on the facts of this case.

324. Use of a sign “in relation to” goods and services means use “for the purposes of distinguishing” the goods and services in question from those of other suppliers, and use such as to create the impression that there is a material link in the course of trade between the goods or services concerned and the undertaking from which those goods or services originate (see *Merck KGaA v Merck Sharp & Dohme Corp* [2017] EWCA Civ 1834 [2018] E.T.M.R.10 (“**Merck**”), per Kitchin LJ at [172]-[173] and the extracts from the decision of the Court of Justice in *Céline SARL v Céline SA* (C-17/06) (“**Céline**”) at [20]-[27] as set out by Kitchin LJ at [174] in *Merck*).

325. It is clear from these passages in *Céline* that use (without more) of a business or trade name will not be in relation to goods and services. This is because:

“21. The purpose of a company, trade or shop name is not, of itself, to distinguish goods or services (see, to that effect, Case C-23/01 *Robelco* [2002] ECR 1-10913, paragraph 34, and *Anheuser-Busch*, paragraph 64). The purpose of a company name is to identify a company, whereas the purpose of a trade name or a shop name is to designate a business which is being carried on. Accordingly, where the use of a company name, trade name or shop name is limited to identifying a company or designating a business which is being carried on, such use cannot be considered as being ‘in relation to goods or services’ within the meaning of Article 5(1) of the directive.

22. Conversely, there is use ‘in relation to goods’ within the meaning of Article 5(1) of the directive where a third party affixes the sign constituting his company name, trade name or shop name to the goods which he markets (see, to that effect,

Arsenal Football Club, paragraph 41, and *Adam Opel*, paragraph 20)’.

23. In addition, even where the sign is not affixed, there is use ‘in relation to goods or services’ within the meaning of that provision where the third party uses that sign in such a way that a link is established between the sign which constitutes the company, trade or shop name of the third party and the goods marketed or the services provided by the third party’.

326. At [275] in *Merck*, Kitchin LJ went on to say that:

“The question must always be whether the activity complained of constituted use of the offending sign in the UK and in such a way that consumers were liable to interpret it as designating the origin of the goods or services in question”.

327. The issue of how the use of the sign will be perceived by the average consumer is a matter for the court to assess on the evidence (Kerly at 16-037).

Identity of Mark and Sign

328. It is not in dispute that:

- i) The perception of identity between sign and mark must be assessed globally from the perspective of the average consumer (*S.A. Société LTJ Diffusion v Sadas Vertbaudet SA* (Case C-291/00) [2003] FSR 34 (“**LTJ Diffusion**”) at [52]),
- ii) A sign will be identical with the registered trade mark “where it reproduces, without any modification or addition, all the elements constituting the trade mark or where, viewed as a whole, it contains differences so insignificant that they may go unnoticed by an average consumer” (*LTJ Diffusion* at [54]). In *Reed Executive Plc v Reed Business Information Ltd* [2004] EWCA Civ 159; [2004] RPC 40 (“**Reed v Reed**”) at [29] Jacob LJ observed that the Advocate General’s opinion was helpful in that case in indicating the sort of difference which would be so minute as to leave the mark and sign ‘identical’ and that which would not:

“...the reproduction of [the plaintiffs’] mark in the same distinctive script but without the dot under the initial ‘A’ might well have been perceived by the average consumer as identical to the original (the change being minute and wholly insignificant) whereas the use of a noticeably different script and/or the addition of another name might be seen as only similar (such changes, at least taken together, being substantial).”

This is borne out by *WebSphere Trade Mark* [2004] EWHC 529 (Ch) [2004] FSR 39 (“**WebSphere**”) in which the presence of a hyphen in the sign (Web-Sphere) was held by Lewison J (as he then was) to be an insignificant difference

from the word mark WEBSPHERE, which would go unnoticed by the average consumer.

- iii) In the case of a word mark, both the visual and aural impression must be considered (*Reed v Reed* at [32] cited by Arnold LJ in *Easygroup Ltd v Nuclei Ltd* [2023] EWCA Civ 1247 (“*Easygroup*”) at [49]).
- iv) There is no reason to suppose that the Court in *LTJ Diffusion* meant to soften the edges of “strict identity” very far, because even where a mark and a sign are not identical, the proprietor of the mark will be protected if there is likelihood of confusion (*Reed v Reed* at [27] and *Easygroup* at [49]).

Identity of goods or services

- 329. Although Getty Images said very little in their opening written submissions about the law on identity of goods or services (referring only to passages in the first instance judgment in *Montres Breguet SA v Samsung Electronics Co Ltd* [2023] EWHC 1127 (Ch), [2023] FSR 1 (“**Montres Breguet [First Instance]**”) which in fact deal with similarity), I did not detect any real dispute over the applicable legal principles in closing.
- 330. Whether a sign is used in relation to identical goods or services will depend on the goods or services listed in the registration of the mark and relied on by the claimant in its pleaded case. Where general terms are used to describe goods or services, they must be interpreted as including goods or services clearly covered by the literal meaning. Goods can be treated as identical where one is a subset of the other (*Montres Breguet [First Instance]* per Falk J at [119]-[120]).
- 331. However, trade mark specifications are concerned with use in trade, so when construing a word used in a trade mark specification the court is concerned with “how the product is, as a practical matter, regarded for the purposes of trade” (*British Sugar Plc v James Robertson & Sons Ltd* [1996] RPC 281 (“**British Sugar**”), per Jacob J at p.289).
- 332. In *Easygroup*, Arnold LJ said that “[w]hether there is identity of services depends in part on the interpretation of the specification of services of the trade mark” and he cited the following passage from *Reed v Reed* with approval:

“...specifications for services should be scrutinised carefully and they should not be given a wide construction covering a vast range of activities. They should be confined to the substance, as it were, the core of the possible meanings attributable to the rather general phrase”.

- 333. This passage was also referred to with approval by the Supreme Court in *Skykick UK Ltd v Sky Ltd* [2024] UKSC 36 (“**Skykick SC**”). Lord Kitchin JSC (with whom the remaining members of the Court agreed) held, *inter alia*, that Cloud Migration is not an “electronic mail service” (and so is not included in “electronic mail services” in class 38 of the Nice Classification) and he agreed with Sir Christopher Floyd and the Court of Appeal that the Judge’s approach of “extending the core meaning of the expression [electronic mail services] to an unclear and indeterminate range of services connected

to electronic mail in an unspecified way” had been contrary to principle. At [365] Lord Kitchin observed that:

“The correct approach, as a matter of principle, in considering a specification of services which is defined by terms which are not clear or precise, is to confine the terms used to the substance or core of their possible meanings: see, for example, *Reed Executive pic v Reed Business Information Ltd* [2004] EWCA Civ 159; [2004] RPC 40, at para 43. So too, if a specification of goods is defined by terms which are ambiguous, then it should be confined to those goods which are clearly covered. These principles are consistent with first, the requirement that the specifications of goods and services must be clear and precise so that others know what they can and cannot do; and secondly, general fairness because any ambiguity is the responsibility of the owner of the mark. If despite this, the words used are still unclear so that they cannot be interpreted, then it is permissible to disregard them. But, in my opinion, that will rarely be the case”.

334. The construction of a specification must be determined as at the date of registration. In *Reed v Reed* Jacob J put the point thus (at [46]-[49]):

“46. ...Mr Howe contends that a specification of goods or services cannot change its meaning with time. Mr Hobbs submits that it can—if the “core” nature of the service changes with time, then the meaning changes too.

47 I have no doubt that Mr Howe is right. One can test the point best by reference to a registration qualified by the words “included in this Class”. From time to time, though rarely, the class in which a particular kind of article is put is changed by international agreement. If that happens it is inconceivable that the trade mark owners’ rights could be changed. “Included in Class X” must mean “included at the time of registration.”

48 So also for a word or phrase which changes its meaning over time. **But that must in practice be very rare. Indeed I know no instance of it in any reported case. The ordinary case—and I think this is one—is where some new variant of an article or service comes into existence after registration. The issue then is whether that new article or service falls within the meaning of the existing specification.** Columbia Graphophone’s Trade Marks (1932) 49 R.P.C. 621 is a good example of a new article falling within an old specification. The specification of goods of the mark under attack was “all goods in Class 8”. Class 8 (of the old classification, bizarrely taken from the classification used for the Great Exhibition) quaintly read “philosophical instruments, scientific instruments and apparatus for useful purposes; instruments and apparatus for teaching”. An application for partial rectification was made on the grounds of

non-use. The excision sought was essentially for “cinematograph films, talking and silent” (my précis). The trade mark owner, a record company, sought to justify the non-user by saying that talkie films had only just become possible and that amounted to special circumstances. The argument failed and excision was ordered. There was an express excision from the specification of a kind of article (talkies) that did not exist at the time of registration” (**emphasis added**).

Application of the law to the facts

Use of the Sign in the course of trade

335. There is a dispute between the parties as to whether *Daimler*, *Google France* and *Coty* assist Stability on the facts of this case or whether, as Getty Images contend, these were exceptional cases, each concerned with very specific factual scenarios which do not apply here. Before I can deal with this issue, I must examine the facts in each case.
336. In *Google France* the CJEU held that, whereas an advertiser might use a sign in the course of trade by bidding on it as a keyword in a keyword advertising service such that it appears in advertisements generated following searches containing that keyword, the service provider (in that case Google) was simply allowing the advertisers to use the signs. Google was not using them by storing as keywords signs identical with trade marks or by organising a display of advertisements on the basis of those key words: “The fact of creating the technical conditions necessary for the use of the sign and being paid for that service does not mean that the party offering the service itself uses the sign” (at [57]). This did not amount to Google’s “own commercial communication”. A similar conclusion was reached in *L’Oréal v eBay*, where the Court held that (in relation to the operation of an e-commerce platform), the use of signs identical or similar to trade marks in offers for sale displayed in an online marketplace is made by the sellers who are customers of the operator of that marketplace and not by the operator itself (at [103]).
337. In *Daimler*, the question that arose for the court was whether the defendant (who had at one time been authorised to, and did, use Daimler’s trade marks to describe itself as an “authorised Mercedes-Benz dealer” in its advertising through a third party online provider) made use of a sign even where (after termination of that arrangement) the defendant had sought to remove all such references from the online advertising. The court held that an advertiser could not be held liable for the acts or omissions of an advertising provider who, intentionally or negligently, disregards the express instructions given by that advertiser who is seeking specifically to prevent that use of the mark. In those circumstances, the publication of that reference on the referencing website can no longer be regarded as use of the mark by the advertiser (at [34]). Similarly, an advertiser cannot be held liable for the independent actions of other economic operators with whom the advertiser has “no direct or indirect dealings and who do not act by order and on behalf of that advertiser, but on their own initiative and in their own name” (at [36]). At [39]-[40], the court identified the broad principle that “using” involves “active behaviour and direct or indirect control of the act constituting the use”, but observed that that will not be the case if that act is carried out by an independent operator without the consent of the advertiser or even against his express will.

338. In *Coty*, a perfume distributor and licensee of the trade mark DAVIDOFF (registered for perfumes) brought proceedings against Amazon for trade mark infringement in respect of Amazon's storing and dispatching of 'Davidoff Hot water' brand perfumes offered for sale by a third party seller via the Amazon website. The question referred to the Court of Justice for a preliminary ruling was whether a person who, on behalf of a third party, stores goods which infringe trade mark rights, without being aware of that infringement, must be regarded as storing those goods in order to offer them or put them on the market, if that person does not itself pursue those aims. The Court reiterated that "using" involves "active behaviour and direct or indirect control of the act constituting the use" and that "only a third party who has direct or indirect control of the act constituting the use is effectively able to stop that use" (at [37]-[38]). At [44], the Court noted that the types of use which a trade mark proprietor may prohibit refers specifically to "the offering of goods, their being put on the market, their being stocked "for those purposes" or the supply of services **under the sign concerned**" (emphasis added). It followed that, in order for the storage of goods bearing the signs complained of to amount to "use" it was necessary "for the economic operator providing the storage itself to pursue the aim referred to by these provisions, which is offering the goods or putting them on the market" (at [45]). Because it was only the third party who intended to offer the goods or put them on the market, their storage by Amazon did not amount to use of the sign in their own commercial communication (at [47]).

339. Getty Images contend that (unlike in *Google France*), Stability is using the sign for its own commercial communication: the communication that bears the watermark* in the form of the output image is the commercial communication of Stability because it is generated by its Model. This, says Getty Images, involves more than merely storing the output images (unlike in *Coty*) – but instead involves offering the service of generating the images and putting those images onto the market. Unlike in *Daimler*, this case also involves active behaviour and control on the part of Stability because (i) it is the entity that trained the Model; (ii) it is the entity that could have filtered out watermarked images in order to ensure that its model did not produce outputs bearing watermarks*; (iii) it makes the Model available to consumers through GitHub, Hugging Face, the Developer Platform and DreamStudio (which I have accepted in relation to v2.x; SD-XL and v1.6. For v1.x the position is more complex); and (iv) it is the entity making the communication that bears the relevant signs. None of this can be said to be the independent action of another economic operator.

340. I agree with these submissions.

341. In oral closing submissions, Mr Cuddigan countered Getty Images' case by submitting that it is the user who is completely in control of the Model. However, upon being questioned about that assertion, he accepted that it was an overstatement and that control is a "nuanced" issue. In my judgment, he was right to do so.

342. Mr Cuddigan put to Professor Farid in cross examination that "the model is a tool controlled by the user and the more detailed the prompt is, the more control is being imposed". Professor Farid responded (and I accept):

"That is partially true. The user has control over what is prompted. However, what the user does not have control over is what the model was trained on. We have no control over that. What the user does not have control over are any semantic

guardrails that might be put on the prompt and any semantic, image guardrails that are put on the output. Absolutely, the user has control over what it asked for but does not have 100% control over what is coming out the other end”.

343. Mr Cuddigan put the point again suggesting that “the user has a high degree of control which is exerted through the specificity of the prompt”, to which Professor Farid responded (and I accept):

“I do not know that I would say that. You cannot separate what comes out of the model from what it was trained on. For example, if the model was never trained on images of dogs, I can ask for images of dogs all day long, but I am not going to get them. I have no control over that

...

I will agree that if -- the more that is specified in the prompt, the more control the user has on what is created. I do not agree, however, that they have complete control. For example, if I open Adobe Photoshop and I paint an image, Photoshop gives me all the tools to do that. I have 100% control because I am dropping every pixel into the image. You cannot say that about Stable Diffusion and these AI models so the more I specify in the prompt, the more control I have, but at some point I push a button and I lose control over what will come out the other end of that model”.

344. The factual evidence added to the overall picture that while users have some degree of control they do not have complete control:

i) Ms Hodesdon was asked in cross examination whether Stability had cause to look at the kind of outputs the Model is producing. She explained that she was responsible for ensuring the correct running of various elements of the inference server, while Stability had a safety team that was concerned with applying filters at the point of inference. She described the managed service provided by Stability to users on its platform (which she later explained was both DreamStudio and the Developer Platform):

“A. In one sense, where we run the models as part of a wider service, we have this thing called an inference server.

Q. Yes.

A. Which will serve outputs from the model and in that we include some pre-processing and some post-processing and the post-processing will look at the output of the model and will filter it against harmful content or other content that we do not want to return to customers and then, based on the output of those filters, will either return a particular code with an error message

or will return the output of the model. This is something we offer in this managed service which we run on our platform”.

ii) Ms Hodesdon was then asked whether Stability might be aware that at various points in time Models v1.x, v2.x (and SD XL) could produce outputs that are NSFW, to which she responded:

“Yes, what I just described with the post-processing was these filters we run as part of our server, which is something we deploy on our infrastructure and it is managed. The model itself is just a sub-component of that much wider pipeline. The model itself can produce outputs. You know, in full generality, it can produce any output. So yes, the model itself can produce outputs that are not safe for work, but where we deploy the model as part of this wider pipeline, we work to, we have checks at the input stage and the output stage to ensure that we are not serving that content back to end users”.

iii) Ms Hodesdon also described the use of a “filtering functionality” that will not render photo-realistic likenesses of well-known public figures, added to the Models to combat concerns over fake news and propaganda. She explained that for prompts provided to the Stability platform, the prompts will be scanned for celebrity names and (where necessary) modified.

iv) In cross examination, Mr Auerhahn confirmed that users pay Stability to access DreamStudio and the Developer Platform. The following exchange then took place:

“Q. They are paying Stability to access and use the software that Stability is providing in order to generate images, are they not?

A. Yes; that is correct.

Q. The software is taking the prompt that the user puts in and the model is then processing it in order to churn out an image.

A. Yes, that is more or less correct.

345. In light of this evidence, I reject Stability’s submission that for each of the access mechanisms referred to above, it is the user and not Stability who is wholly responsible for conditioning the circumstances in which the outputs are generated. While the evidence shows that Stability exercises a greater degree of control over the outputs generated by users accessing via DreamStudio or the Developer Platform than it does in respect of users who have downloaded the Model (who are able to modify the code and run additional training with their own data), nevertheless:

i) Stability has made v1.x available via the API/Developer Platform and v1.4 available via DreamStudio, thereby taking responsibility for those releases (including for the data on which relevant versions were trained) and employing a post-processing filtering service;

ii) Stability has trained v2.x (and SD XL and v1.6) from scratch. From v2.x onwards, Stability has made deliberate choices as to the content and make-up of the dataset on which the Models were trained and the filters (if any) to be applied at that stage.

346. Furthermore, Stability is responsible for the model weights for the versions referred to above. Also, as Mr Cuddigan confirmed in closing, Stability has written the code that users will download if they are accessing these Models (from v2.x onwards) via GitHub or Hugging Face and which they will then run on their local computers.

347. To my mind, this goes beyond merely creating the technical conditions necessary for the use of the Sign. The provision of access for users to the Models via the different access mechanisms cannot be equated merely with the storage of goods by Amazon (*Coty*), or with Google allowing advertisers to use the signs (*Google France*), or with eBay making its online marketplace available to customers (*L'Oréal*).

348. While an AI model such as Stable Diffusion may be described (in simple terms) as a tool to enable users to generate images, that is not a complete description. As the Experts agreed in the Technical Primer, Stable Diffusion is a machine learning system which derives its primary function largely from learning patterns from a curated training dataset. Its final function is not directly controlled in its entirety by the engineers who designed it, but a large part of its functionality is indirectly controlled via the training data. The model weights are learned from the training data and it is the model weights which control the functionality of the network. Although the process of inference does not require the use of training data, the outputs generated during inference will (at least indirectly) be a function of that training data. Thus, as the Experts agree, the generation of watermarks* by the Model “is due to the fact that the model was trained on some number of images containing this visible watermark”. This is the responsibility of Stability. Under cross examination, the following exchange took place between Ms Lane and Professor Brox:

“Q. The reason why all these different prompts produce watermarks is because, presumably, Stability failed to remove images bearing Getty Images watermarks from the dataset it trained on?

A. Yes; that is my guess”.

349. It is Stability’s case, with which Getty Images agree, that “[n]o relevant consumer will seek to generate or use content containing watermarks* as such watermarks mar the appearance of images and render them unusable for most purposes”. In such circumstances, I agree with Getty Images that it is very difficult to see how it can sensibly be argued that the user of Stable Diffusion (by whatever access mechanism) has any control over when a watermark* is produced. As Getty Images submit, the only entity with any control in any meaningful sense of the word over the generation of watermarks* on synthetic images is Stability. It is certainly not “passive” as Stability submits.

350. Even if (consistent with Stability’s submissions on context but inconsistent with my findings above) the average consumer will read and remember in detail the terms and conditions for use of the Model and/or the License terms and will see and understand

that they are said to be wholly responsible for their own synthetic output images, that does not mean either that Stability in fact has no control over the images that are generated, or that it has deployed no active behaviour in relation to the generation of those images. I consider Mr Cuddigan's submission in closing to the effect that "all the actions are [those of] the user" to be unsustainable in light of the evidence.

351. Because I detected in Stability's submissions a considerable degree of overlap between its arguments on "use in the course of trade" and "use in relation to goods and services", I shall not try to summarise my conclusions on use in the course of trade at this juncture but will instead turn now to consider the points that I understood Stability to be making under both heads. I shall return to my overall conclusions on use at the end of this analysis.

Use of the sign "in relation to goods and services"

352. Stability submits that there is nothing in the evidence to support a finding that use of the Sign by Stability is in the context of a commercial activity or that it is "in relation to" goods and services. Stability relies heavily upon the fact sensitive and contextual nature of the enquiry under this head, which must be assessed having regard to the perception of the average consumer. It refers me to *Trebور Bassett v the FA* [1997] FSR 211 ("**Trebور Bassett**"), in which Rattee J rejected the argument that publishing and marketing cards showing photographs of players wearing the England team football strip (including the three lions logo) was use of that logo or indeed, use of that logo as a sign in respect of the cards. Essentially it contends that the appearance of watermarks* on synthetic images is not "trade mark use" in that it is not liable to affect the functions of the trade mark, in particular its essential function of guaranteeing to consumers the origin of the goods (see C-206/01 *Arsenal v Reed* (ECJ) [2003] Ch 454 at [51]).

353. Stability contends that the average consumer seeing the Sign on a synthetic image will not form the impression that it is a message about trade origin or that it indicates any material link in the course of trade between the images on which the watermarks* appear and Getty Images.

354. Getty Images submit, on the other hand, that the Signs would plainly be perceived as being used in relation to goods and services given that they are affixed to synthetic image outputs and they generally appear in the same location on the synthetic images as the Marks appear on Getty Images assets, i.e. they are positioned within a grey translucent banner which is overlaid on the synthetic image in the bottom third on the right hand side. Getty Images say that use of the Sign in this manner indicates to consumers that the synthetic image (i) originates from Getty Images – for example because consumers are very accustomed to the use of watermarks on photographs, such as children's school photographs, in everyday life; or (ii) is in some way connected to Getty Images, for example that it was generated by a Model that was trained on Getty Images' assets under license from Getty Images.

355. To address these opposing submissions, I need to consider in detail (as both sides did in closing) the four grounds identified by Stability in opposition to Getty Images' case.

(i) Responsibility

356. First, Stability contends that consumers understand that they are responsible for the images generated in response to their prompts and it relies upon its submissions as to context. There is a considerable overlap here with Stability's arguments on "control". I have already held that the average consumer who downloads the Model for use locally will understand (whether they have read the terms and conditions of the License or not) that the images they generate are generated in response to their own bespoke prompts. They will not think that Stability is directly dictating the images generated by the Model. I have also held that users accessing the Model via the API and DreamStudio will understand that the images have been generated by an AI model in response to their own prompts.

357. However, I reject the suggestion that the average consumer (using any of the three access mechanisms and whether he or she has read the terms and conditions, the Model Card or the License, or not) will regard the output of the Model (at least in so far as that output bears watermarks*) as solely their responsibility. Indeed it seems to me that Getty Images are right to say that this proposition is contradicted by Stability's own case which contends (as set out above) that consumers will not wish to generate watermarks*. The average consumer who does not wish to generate a watermark* and who has not entered a text prompt designed to do so is, in my judgment, most unlikely to think that the generation of such a watermark* is solely his or her responsibility whatever he or she may have read in the terms and conditions of use or elsewhere.

358. In closing, Stability sought to rely upon submissions made by Getty Images' counsel to the Court of Appeal following my decision to exclude references to CSAM from Getty Images' case at trial, as follows:

"Getty is not alleging, or seeking to allege, that Stability is criminally responsible for what a minority of its users, one might say perverted users, may do with Diffusion. **The offence is committed by those users who prompt the model to produce CSAM, just in the same way as it is they, rather than Stability, who may produce, may choose what prompts they use and thereby generate, if we are right, other forms of pornography and violent images and images of the nature of propaganda.** Of course, from Getty's perspective, if and when the offending images are put in circulation, the damage is the same..." **(emphasis added).**

359. Stability contends that these submissions recognise that criminal responsibility for the production of CSAM images lies with the user and not with Stability and that it is users who choose their prompts and "thereby generate" images. This is true, but it is not the full story. To my mind it does not provide support for the submission that Stability then goes on to make to the effect that "even for benign prompts which produce watermarks* it is not easy to see why the same regime of responsibility would not apply". One very obvious distinction is of course that, unlike pornography or violent images, watermarks* may appear where the user does not seek to generate them.

360. Thus the perception of the average consumer as to 'responsibility' for the output image may, in my judgment, depend on whether he or she has deliberately sought to generate it. Stability itself argues in its closing submissions, "when a consumer knows they are committing misuse, or making out of scope use, they will surely not consider Stability

is responsible for the resulting images, or as making any commercial communication through those images”. Although this proposition begs the question as to what the average consumer will ‘know’ (a question which goes to context), I note that it is not suggested by Stability that a watermark* will only be produced in response to a prompt when the average consumer is misusing the Model or engaging in out of scope use. Why should the average consumer consider himself responsible for the appearance of such a watermark* when it appears unprompted and apparently randomly?

361. Stability contends that its point on out of scope use applies to users adopting the words “news photo” to prompt output images from the Model. This is because (having regard to Stability’s submissions on context) it says that the average consumer using such a prompt will understand it to be out of scope (i.e. designed to produce “factual or true representations of people or events” or, in other words a simulacrum of a photographic image). While, on balance, I consider that the average consumer who has downloaded the Model will understand from a reading of the Model Cards and the GitHub and Hugging Face pages that such use is out of scope of the Model, I reject the suggestion that this understanding will lead them to regard themselves as “responsible” for the production of a watermark* on an out of scope image or will affect their perception of any connection between the synthetic image and Getty Images. On balance they will understand that the production of the watermark* is a function of the way in which the Model was trained – for which they were not responsible. Further and in any event, out of scope use is not misuse.
362. Further, I disagree with Stability that the average consumer using DreamStudio or the Developer Platform will understand that prompting the Model with the words “news photo” would be ‘out of scope’ for the Model. Some may, but a significant proportion will not. In any event, I do not consider that the average consumer (whatever his level of understanding) would regard himself as responsible for the generation of an unexpected and unwanted watermark*. I do not consider it enough that the average consumer will understand the Model to be designed to create artificial images.

(ii) Technical

363. Second, Stability contends that the average consumer has a sufficient technical understanding to appreciate why the watermarks* appear. It says that this understanding is in itself enough to dispel any perception of an origin message from Stability.
364. I have already expressed the view that a significant proportion of sophisticated users accessing the Model by downloading it onto a local computer or via the Developer Platform will understand why the watermarks* appear. This is borne out by observations in the Reddit Exchanges (which, are predominantly made by obviously sophisticated consumers):
 - i) “Looks like Stable Diffusion was trained on watermarked images – when asked for vector art, it put the iStockPhoto watermark all over it”;
 - ii) “It’s essentially trained on the internet. If you pull on something that mainly has watermarks. It will gen watermarks”;

- iii) “It’s not a watermarked image. The AI generated the watermark because it saw the watermark on enough images in its training data that it has associated the watermark with stock photos”;
- iv) “There’s AI that can remove watermarks. They’ll need to unleash that on the dataset for future training”;
- v) “I swear one day I’m going to get a middle finger image with ‘TRAIN ON THIS YOU A****LES’”;
- vi) “...The model is just so used to seeing Getty’s watermark on many images that it thinks that’s an appropriate element to generate sometimes, because it has learned that as a common image trait, the same way it has learned any other concept”;
- vii) “Literally no one is disputing it trained on Getty Images”;
- viii) “It is probable that SD trained off of Getty images...”;
- ix) “if the model sees something enough it replicates it, watermarks tend to do that”.

365. The August 2022 SalesForce Message, the Japanese Garden Post and a post on Mastodon dated 28 December 2022 attaching an image of the Northern Lights (“**the Northern Lights Post**”) which is included in Annex 8I, all evidence a similar understanding. The Northern Lights Post records that:

“It seems that #StableDiffusion overfits a bit when prompted with northern lights. I’m trying to get it to create a drawing instead of a photo and all I get are nighttime (sic) photos with GettyImages watermarks.”

366. I consider that the fact that these (sophisticated) average consumers understand how AI models work and why watermarks* are generated is likely to preclude any possibility that they will perceive the synthetic image to which the watermark* is affixed to originate directly from Getty. However, I cannot see why this understanding precludes a perception that the provision of the synthetic image is in some way connected to Getty Images, perhaps because the Model has been trained on images licensed for use by Getty Images. The observations on Reddit to which I have referred above evidence a perception that the Model has been trained using Getty Images Content. Some appear to evidence an understanding that this will have been without the consent of Getty Images, but others do not. The October 2022 Salesforce Message is evidence of a user who plainly thinks that licensing by Getty Images may have occurred.

367. As for the average consumer using the DreamStudio mechanism of access, again I have already expressed the view that his or her level of technical expertise will be less. Some will understand why watermarks* are generated (see by way of example the user who emailed DreamStudio in February 2023 about the fact he was generating watermarks* - in response to the DreamStudio Email he says: “Thanks for your answer but i know this behaviour and perfectly understand why it happen...”), but a significant proportion will not.

368. In such circumstances, I cannot see that this argument as to “technical” understanding can apply to them. In my judgment, a significant proportion of this class of average consumer will perceive the presence of a watermark* on the synthetic image to indicate that it originates from Getty Images or that the provision of the synthetic image is in some way connected to Getty Images, perhaps because Getty Images has licensed Stability to train its Models on Getty Images Content.

369. One good example of a lack of technical understanding on the part of a user appears in the Reddit Exchanges:

“Soo I had a surprise in one of the generations. i know many would misconstrue this as evidence of ai stealing or whatever which it isn’t but why does it generate accurate watermarks like this?? Does this mean it can generate other copyright protected stuff unintendedly??”;

370. That the buttons and branding on the DreamStudio website inform DreamStudio users that they are receiving a generated image using Stable Diffusion does not appear to me alter the position. Stability relies on evidence given by Mr Stanley about images appearing on the Getty Images Website, shown to him in cross examination. Specifically, Mr Stanley was shown a screenshot of the Getty Images Website with a favicon  at the top of the page and the website name also appearing: [gettyimages.co.uk/search/2/image?family=editorial&page=5&phrase=ces%20expo&sort=best](https://www.gettyimages.co.uk/search/2/image?family=editorial&page=5&phrase=ces%20expo&sort=best). The title of the page was “ces expo” and it showed nine “editorial images” under that heading, including images of open laptops displaying brand names such as Dell and Lenovo and an image of a sign displaying the words “intel ai”. Mr Stanley agreed that this image would not be understood by a user looking at this page as suggesting that they were on an Intel website or that they were dealing with Intel – instead it would merely indicate to the user that an editorial photographer based with Getty Images had taken a photo of an Intel logo.

371. I do not consider Mr Stanley’s evidence on this point, however, to bear the weight placed on it by Stability. The web page shows various images taken in connection with the “ces expo” and the context of the appearance of the intel ai sign is thus very different from the context in which the watermark* appears on generated synthetic images. As Stability correctly submits, whether a sign is understood to be a message about trade origin is fact sensitive and contextual.

372. Equally, I do not consider the fact that users will see “a consistent fudge” in respect of the name of the photographer under the watermark* to be sufficient to dispel the perception of brand message from the appearance of that watermark*, as Stability contends. As a point of principle this appears to me to be a somewhat surprising proposition and it is not, in any event, borne out by the evidence.

373. None of the Reddit Exchanges suggest that any real life user discussing the phenomenon of the appearance of watermarks* was focussing on, or was interested in the significance of, the appearance or legibility of the photographer’s name.

374. It is true that Mr Stanley confirmed that the photographer’s name always appears on the Getty Images watermark and it is certainly part of the context in which the Sign will be seen. But I consider that the average consumer (i) may very well not remember that

the photographer's name always appears together with the Getty Images watermark; (ii) is much more likely to be focussing on the Sign itself; and (iii) may think nothing of the appearance of an indecipherable splodge. No doubt if they do remember that the photographer's name generally appears they would, as Mr Stanley said, expect it to be legible. But, as he also said, he could not guarantee that a customer would be surprised if the name was not legible:

"I think they may be surprised, but in my experience, my Lady, there are an awful lot of things that our customers, who are busy people, do not see or notice....".

375. Accordingly I do not consider that the indecipherable nature of the photographer's name will be perceived by the average consumer paying a moderate degree of attention as dispelling the commercial messaging conveyed by the watermark*. It is part of the context in which the Sign is perceived but I do not consider it to be the "red flag" suggested by Stability.
376. Stability made no similar submission in relation to the iStock watermarks*.

(iii) Brand Messaging

377. Third, Stability contends that consumers have certain expectations about how brand messaging is deployed. It says that when dealing with consumer goods, such as clothing, consumers have expectations as to how the origin of the clothing will be designated (the label on the back of the neck; the crest on the chest etc.). Similar principles apply (it says) in respect of use of online services, where consumers will rely on web addresses and page titles and would expect a collaboration between service providers to be clearly explained.
378. I understood Stability to make four main points in support of this contention:
 - i) That the use of brand messaging on the Stability website (which would be seen by the average consumer using DreamStudio and the Developer Platform) would reinforce the understanding of the user that he or she was using a Model provided by Stability.
 - ii) In turn, this would serve to accentuate the unusual manner in which the watermarks* appear. The distorted nature of the watermark* on the output (when combined with the strong branding message) would speak to the proposition that the watermarks* are artefacts of the image synthesis and not some commercial origin message. Even if the watermark* appears clearly in the image, the whole sign is recognisably "off".
 - iii) The average consumer would realise that if a link with Getty Images were the intended message, some other orthodox branding would be present outside the four corners of the image. Even if he or she thought that a watermark might be used for such messaging, they would expect it to be rendered precisely.
 - iv) It is apparent from the contents of the synthetic images themselves that "these are not the sort of images which Getty promulgates".

379. I have already accepted that the DreamStudio user is repeatedly presented with a series of trade marks which indicate that the origin of the service provided is Stability. While this also appears to be the case for the user of the Developer Platform when he or she originally sets up an account on the Stability website, there is no evidence to suggest that their experience is thereafter the same as that of the DreamStudio user.

380. Focussing, however, on users of DreamStudio, I have already found that they would understand that they are on a Stability website when generating images using the Model. But I do not consider that the existence of “conventional branding practices” (in the form of a favicon, a page title and a web address) on that website precludes the potential for the average consumer (a significant proportion of whom will not appreciate how the Model has been trained or what images it may have been trained on) to think that a watermark* appearing on a synthetic image indicates some connection or material link with Getty Images. Nothing that Mr Stanley said about the understanding of users interacting with the Getty Images Website or the use of additional branding (including branding indicative of a collaboration) alters this conclusion.

381. Although it is true that most of the images bearing Getty Images watermarks* in Annex 8H show images which are recognisably “off” (because they are seriously distorted, include more than one image of the same individual, show individuals with unusual proportions or with missing or additional limbs or distorted features and the like) that is not the case for the “real world” images in Annex 8I which bear the Getty Images watermark* and in respect of which the versions of the Model which generated them are known: for example, the Musician Image, the Hugging Girls Image (which merely looks to be improperly centred) and the two Japanese Temple Garden Images. I do not consider that the average consumer would consider these to be “recognisably off” and nor do I think that the average consumer would regard these as images of the sort they would not expect to see produced by Getty Images (which uses photographers to create conceptual images as well as to capture photographs of real events and real people). Mr Barwick observed in his statement in very general terms that he had “not seen anything that [he] would confuse with real images”, but Mr Barwick is a highly skilled professional photographer. With the utmost respect to him, I do not consider that his views can be equated with those of the average consumer who, as I have found, is paying a moderate degree of attention.

382. I set out again two of the Annex 8I images below, namely the Musician Image and the First Japanese Temple Garden Image:



383. I also bear in mind that the synthetic images produced by the Model appear to display the Getty Images watermark* in the conventional way that watermarks are used on photographs, and (at least in some cases) in the same general position. It is true that the watermarks* in the above images include some additional elements: the watermark on the Musician Image has an “m” with three arches (and thus four legs)); while the First Japanese Temple Garden Image watermark* appears to include an extra letter or character between the “a” and “g” and some letters are distorted. Furthermore, the photographer’s name in each image is heavily distorted. But, while these features might alert some people to the fact that the watermarks* are not, as Stability puts it, “the real thing”, I consider that there would be a significant proportion of average consumers who would not be so alerted.

384. While I accept Mr Stanley’s evidence that one would expect the Getty Images logo to be precisely rendered on the Getty Images Website, I do not consider that the rendering of the watermarks* in these images is so imprecise as to displace the impression that these images originate from Getty Images (at least for users of DreamStudio) or that the service that is being provided by Stability is in some way connected to Getty Images, perhaps because Getty Images has licensed its Content for use by Stability in training the Model (for all users), or perhaps because the images are in some way licensed to Getty Images.

385. Of course, as I have already said, it is a novel feature of this case that every time a watermark* is generated it will be generated in relation to a unique image and will or may appear in different forms. This is because the Model learns from its training data the statistics of patterns or probability distribution associated with certain concepts – the Model then generates new images by sampling from the distribution. Thus there are many ways in which one synthetically generated Getty Images or iStock watermark* will appear visually different from another synthetically generated Getty Images or iStock watermark*. That makes this a highly fact sensitive enquiry, which inevitably requires consideration of both the specific image and watermark* generated. The more distorted, blurred or recognisably “off” the image and/or the watermark* is, the more likely it is that the average consumer will realise that there can be no material link or connection of any kind with Getty Images.

386. In closing Stability relied upon an observation from Ms Cameron to the effect that an image generated by one of the users posting as part of the Reddit Exchanges was no

more than “very reminiscent” of the Getty Images watermark. This image illustrates the point I have made in the previous paragraph nicely:



387. Aside from the fact that the relevant Model used to generate this image (which appears in Annex 8I) is not known, this is an example of a watermark* which is very obviously imprecisely rendered, unclear and noticeably “off”, including by reason of the nature of the image on which it is rendered. But it is not representative of all watermarks* generated in real life, as the Musician Image and the Japanese Temple Garden Images show. Watermarks* could seemingly be rendered in a wide variety of different forms on an infinite variety of different images.
388. As for the available images bearing the iStock watermark* - these show designs which do not purport to be in any way photo-realistic: see the Spaceship Image and the Dream Image set out later in this judgment at paragraph 402. Accordingly, I fail to see how Stability’s submissions about the content of the synthetic images themselves could possibly apply to these. Some of the watermarks* on the iStock images in Annex 8H and Annex 8I are slightly distorted and blurred, but others, including the two images on which Getty Images relies for the comparison between Mark and Sign are much less so and I do not consider the average consumer looking at the Spaceship Image and the Dream Image would consider them to be recognisably “off” or to speak to the proposition that the watermarks* are purely the artefacts of the image synthesis process rather than the purveyors of some commercial origin message.
389. In response to Getty Images’ point that the average consumer may think that a synthetic image bearing a watermark* was trained on Getty Image assets under license by Getty Images, Stability makes a couple of points. First it submits that watermarks* are “on any view a rare occurrence” and that the average user of any of the access mechanisms will have generated a host of synthetic images before ever coming across one. Thus, it says, it makes no sense for a particular image to indicate some further licence arrangement. I reject this argument. I cannot say whether the generation of watermarks* by v1.x or v2.x is a rare occurrence for reasons I have explained when dealing with the threshold question. Some of the observations in the Reddit exchanges might suggest otherwise. Further there is no basis whatever for me to determine that the average consumer seeing a watermark* will already have generated numerous

images and will not therefore perceive that there may be a licence arrangement with Getty Images (at least in respect of some images).

390. Second, Stability contends that the average consumer (who I accept will be used to seeing legitimate watermarks appearing on a variety of images) will “of course be familiar with the reason why watermarks exist: it is to render an image commercially unusable unless a licence fee is paid”. This is consistent with Mr Stanley’s evidence that one reason for including watermarks on Getty Images Content is “to stop unauthorised use of imagery”. However, I do not see how this assists Stability. The user’s familiarity with the use of watermarks and the rationale for that use seems to me only to enhance the perception (at least where the unsophisticated average consumer using DreamStudio is concerned) that a watermark* appearing on a synthetic image is intended to convey a message that the image should not be used because it is in some way licensed to Getty Images.

(iv) Lack of Evidence

391. Finally, Stability contends that there is no evidence to support the two distinct propositions advanced by Getty Images in support of its case that the Signs would be perceived by the average consumer as being used in relation to goods and services, namely (i) that consumers will think that the watermarked* images originate from Getty Images; and (ii) that consumers will think that the watermarked* images were generated by a model which was trained on Getty Images assets under licence from Getty Images.

392. Stability says that if this were true, Getty’s Salesforce Materials would be riddled with queries (including about licensing) and it suggests that there is not a single Getty Images customer raising these issues. There are certainly no witnesses before the court to explain that this is how they responded to the generation of a watermark* on a synthetic image output from one of the Models.

393. I reject Stability’s submission (which appears to me to be inconsistent with its case that users will not want to generate images bearing watermarks*) that because a watermark conventionally indicates that a licence is required, Getty Images could be expected to be receiving numerous enquiries about licenses. Users of DreamStudio (when using the Model in default mode) can choose from one of four images on screen; any user (accessing via any of the access mechanisms) can prompt the Model to produce further images. They can play around with the text used in their prompts to achieve the image they want. Because those users will not want to generate images bearing watermarks*, it seems to me that the vast majority are likely merely to discard such images; they are most unlikely to approach Getty Images to ask for, or check, the licensing position. This goes some considerable way towards explaining the lack of evidence of any queries in this case.

394. In any event, however, while it is true that there is very little evidence of queries being raised, the cupboard is not entirely bare. There is in fact the October 2022 SalesForce Message (to which I have already referred) which expressly shows a consumer asking the question:

“...recently I came across a prompt that results in well over 50% of the images showing an iStock watermark, are these images

legal to use in my projects? How does Stable Diffusion go about licensing material from you?

395. This user appears to have formed the impression (i) that Stability had obtained a licence from Getty Images to use these images; and (ii) that the presence of the watermark* on the images might be read as prohibiting his use of them in his projects. This latter understanding is consistent with Ms Cameron's evidence that the presence of a watermark "would certainly make me question if it was licensed or not". Stability points out that following the explanation from Getty Images by way of reply to this message to the effect that there is no affiliation with Getty Images, the user replies "I kind of assumed as much, but if they aren't licensing content from you then please explain the attached Images". Stability suggests that it is clear from this that the user had in fact made no origin association assumption in response to the watermark*.
396. On balance I disagree. This user was sufficiently confused as to origin to cause him to write to Getty Images and pose the questions set out in his email, including the question about whether he could legally use the images he had generated which included the watermark*. I do not consider his eventual reply to establish that there had been no origin association – indeed he still poses a question about how the images could have come about if they are not being licensed by Getty Images.
397. The user raising this issue is based in the United States, but for reasons I have already given, I consider this message to be evidence that Getty Images' case on use of the Sign in the United Kingdom is consistent with reality. Indeed this accords with my own impression. I reject the suggestion that this user is merely an "outlier", unrepresentative of the average consumer. I have already said that various of the Reddit Exchanges also suggest users perceiving brand messaging when encountering a watermark*: e.g. "I can clearly see istock as watermark, not sure istock will approve those generated pictures" and making an association between Stability and Getty Images in respect of the training of the Model.

Conclusion on Use

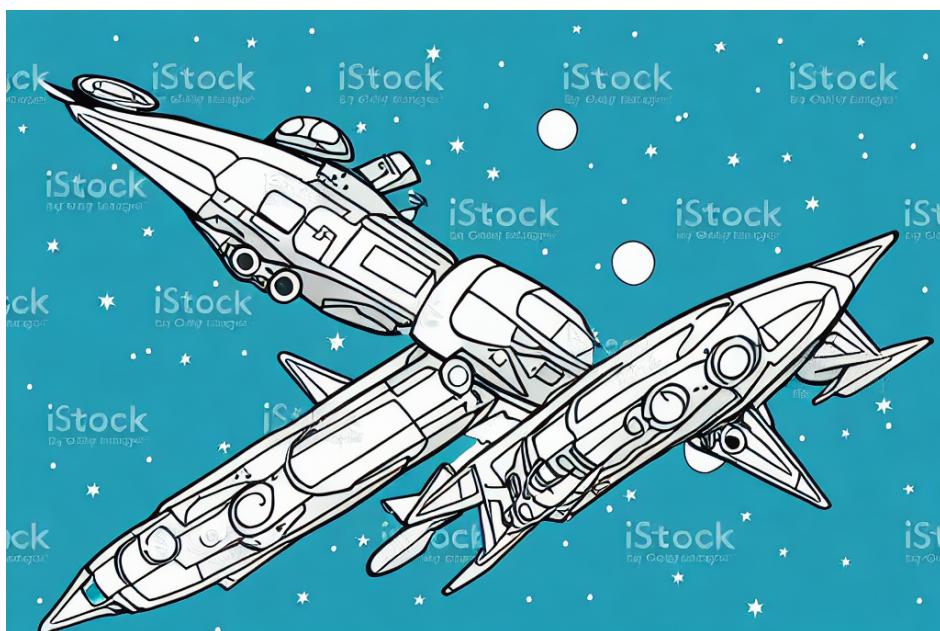
398. For all the reasons I have identified above, I consider that (subject to the points I have made above about the fact sensitive nature of the enquiry in relation to every example of a generated watermark* on a synthetic image) there is evidence in this case of output images, generated by v1.x and v2.1, which include the Sign and which, in my judgment, will be perceived by the average consumer as a commercial communication by Stability. Stability is running a business in the UK and providing Stable Diffusion to consumers as part of that business. The Signs are affixed to synthetic images generated by customers owing to the functionality of the Model, itself dependent upon its training data (over which Stability has absolute control and/or responsibility). It is in this way that Stability "offers and puts synthetic images bearing the Signs on the market" and this is Stability's commercial communication to the consumer. I do not consider this to be in any way analogous to the *Trebor Bassett* scenario, or indeed to the decision in *Merck*, at [275], to the effect that the use of the word Merck "in a context which consumers would understand to be a description of an entity engaged in an activity other than the provision of goods and services" was not trade mark use.
399. For all the reasons I have given, I find that the use of the Sign (assuming it to be sufficiently clear and subject to the fact sensitive analysis referred to at paragraphs 381

and 385 above in relation to the synthetic image on which it appears) is such as to create the impression that there is a material link in the course of trade between the goods concerned and the trade mark proprietor. The average consumer using DreamStudio may interpret the Sign as designating Getty Images as the undertaking of origin of the images. The average consumer using all access mechanisms may interpret the Sign as indicative of a connection between Getty Images and Stability, including because the Models have been trained on Getty Images Content and licensed for use.

400. In my judgment, this use is in the context of a commercial activity with a view to economic advantage (as Mr Auerhahn's evidence confirmed) and is not a private matter. Stability maintains (albeit only in a footnote in closing) that the use in issue in these proceedings is a private use by a consumer. But this submission can only be premised on an assumption of total control on the part of the consumer, and that assumption is, in my judgment, erroneous for reasons I have already explained.
401. Stability made an additional point in its written submissions to the effect that the paradigm trade mark use takes place before any commercial transaction but that the watermarks* in this case will only be seen by consumers long after they have made the relevant commercial choice, namely to download, or sign up to gain access to, Stable Diffusion. Thus, submits Stability, the watermarks* post-date any conventional brand usage. In oral closings, Mr Cuddigan emphasised that this means that Stability cannot be regarded in any real sense as putting images onto the market "under the Sign". However, I do not see how this submission can sensibly be pursued, in light of the decision in *Iconix*. As I also understood Mr Cuddigan to accept, it is a question of fact whether the consumer will perceive some disjunctive branding on an image appearing on a website he or she has already entered as a brand origin message. I have determined that question in Getty Images' favour for all the reasons set out in this section.

Identity of Mark and Sign

402. Pursuant to the Order of 22 May 2025, the comparison between Mark and Sign is to be tried by reference to the following four examples set out below which I shall refer to respectively as "the Obama Image", "the First Japanese Temple Garden Image", "the Spaceships Image" and "the Dreaming Image". The Obama Image and Spaceships Image were generated as a result of the Getty Watermark Experiments and are included in Annex 8H. The Obama Image was generated by v2.1 in response to a verbatim prompt, while the Spaceships Image was generated by v1.2 in response to a "vector art" prompt. The First Japanese Temple Garden Image (generated by v2.1) and the Dreaming Image (in respect of which the relevant Model is not known) were generated by real world users of Stable Diffusion and are included in Annex 8I. The Dreaming Image was generated by a prompt that included the words "vector art".





403. In Opening, Getty Images provided a useful table in which they conducted a comparison of Marks and Sign, which I reproduce below (without Getty Images' accompanying submissions). For the purpose of this exercise I understand Getty Images to have simply extracted the watermark* from the images identified above and inserted it into the table under the heading "Sign". I have added reference to the relevant image in the table.

Trade Mark	Sign	
GETTY IMAGES (UK859, UK005) gettyimages (UK925)	 (Annex 8H/p.84 [A2/1(u)/375])	The Obama Image
	 (Annex 8I/p.24 [A2/1(v)/418])	The First Japanese Temple Garden Image
ISTOCK (UK297, UK819)	 (Annex 8H/p.22 [A2/1(u)/313])	The Spaceships Image
	 (Annex 8I/p.23 [A2/1(v)/417])	The Dreaming Image

404. I have already indicated that UK859, UK005, UK297 and UK819 are all word marks and it is common ground that these registrations cover the word in any typeface/script (*Dreamersclub Ltd's Trade Mark Application* [2019] RPC 16 per Phillip Johnson sitting as an Appointed Person at ([11]-[12])). UK925 is a figurative mark registered in

black and white. Again it is not in dispute that it covers the word in that script in any colour (*Kerly* at [10-093]).

405. In closing, Mr Cronan for Stability submitted that blurring or changes to legibility or the addition of letters is not merely “typeface” or “script”, but amounts to additional visual content. In general terms I accept this submission, but I must consider the comparison to see whether any such additional visual content in fact appears in the example Signs in this case. Mr Cronan also identified a difference between the analysis of the word mark (which shows GETTY IMAGES as two words) and the analysis of the figurative mark (which is all lower case, is written as one word “**gettyimages**” and includes an emboldened part). This was not a difference that Getty Images focussed on in its submissions.

*IStock Watermarks**

406. Looking first at the ISTOCK Marks and Signs, I agree with Getty Images that there is clear identity in every respect: aurally, visually and conceptually. I consider that the average consumer would regard these as identical. I did not understand Stability seriously to suggest otherwise. In its written submissions it made a point about the watermark* in the Spaceships Image being an integral part of the image because it is layered in between the background and foreground content. However, I do not consider this point to be at all persuasive. The Spaceships Image includes a pattern of what appear to be repeated watermarks*, some of which are obscured by the spaceships (as in the partial image included in Stability’s skeleton argument) and some are not. I do not consider the fact that some of the watermarks* are obscured, or layered between the background and foreground is likely to affect the perception of the average consumer of identicality. There is nothing unrepresentative or unrealistic in making a comparison by reference to the Spaceships Image, notwithstanding that it was generated pursuant to the Getty Watermarks Experiments.

407. I have found that users of v1.x of the Model in the real world have used the prompt “vector art” and have generated a synthetic image with an iStock watermark*. Although it is not possible to say to which version of the Model the Dreaming Image relates, it was produced by a user in real life and therefore (for present purposes) appears to me to be representative of a watermark* that could be created using a prompt which includes the words “vector art”.

*Getty Images Watermarks**

408. Turning then to the Getty Images’ Marks and Signs, I need to deal with the two Signs separately.

409. Stability contends that because the watermark* on the Obama Image was generated in the course of the Getty Watermark Experiments, I should disregard it as unrealistic and unrepresentative of a watermark* that would be produced in real life. For all the reasons I have already given, I agree. To my mind there is no evidence that such a clear Getty Images watermark* has been produced by any of the Models in issue using a real world prompt and accordingly I do not consider that it is appropriate to use this Sign for the purposes of the comparison.

410. However, if I am wrong about this, then I agree with Getty Images that the Sign on the Obama Image has clear identity, aurally, visually and conceptually with the Getty Images word mark. There are only two differences between Mark and Sign:

- i) First, the Sign has an apostrophe after the letter ‘s’ which is not present in the Mark. However, I accept that this additional element makes no difference to the aural pronunciation or the conceptual meaning and that (by analogy with the hyphen in *Websphere*) it is a difference so small that it would go unnoticed by the average consumer.
- ii) Second, the Mark has a space between the words “GETTY” and “IMAGES” whereas the Sign does not. Again, this makes no difference to the aural pronunciation or the conceptual meaning and is a difference so small that it would go unnoticed by the average consumer.

411. I arrive at a similar conclusion in relation to the comparison between the figurative Mark and Sign. Here the only difference is the apostrophe and, as I have said, that would go unnoticed by the average consumer.

412. I did not understand Stability seriously to dispute this analysis. It points out that the photographer’s name under the Sign is garbled and that (having regard to the evidence of Mr Stanley) the average consumer would expect the photographer’s name to be legible. It also contends that because the garbled name is co-located with the other wording within the boundary of the watermark*, the average consumer will perceive it as part of the Sign and will not therefore regard the Sign as identical to the Marks.

413. I reject these submissions. I do not consider that the average consumer would perceive the photographer’s name as part of the Sign (as opposed to merely the context in which it appears) and none of the evidence from real life users suggests the contrary. In such circumstances the fact that the photographer’s name is garbled (while potentially relevant to context and addressed above) is irrelevant to the assessment of identity between Mark and Sign.

414. The clearest Getty Images watermark* produced in real life is on the First Japanese Temple Garden Image. Getty Images contends that, once again, (for both the figurative and the word Mark) there is clear identity in every respect: aurally, visually and conceptually. While it accepts that there is a difference because of “the presence of a shadow in between ‘a’ and ‘g’” in the word “images”, it nonetheless contends that this shadow is “hardly discernible and may form part of the underlying image”.

415. Looking at the watermark* as it appears on my screen (i.e. in its proper context on the image), I disagree. As Stability submits, the watermark* on the First Japanese Temple Garden Image appears to have an extra ‘i’ in the word ‘images’ so that it appears to read “imaiges”. The extra ‘i’ is no less “shadowy” than the other letters in the word and it is plainly not a part of the underlying image. It is not merely a difference in typeface or script. On balance I consider that this difference is not so insignificant that it may go unnoticed by an average consumer exercising a moderate degree of attention. The presence of the additional ‘i’ affects the aural analysis, just as it affects the visual analysis.

416. In closing, Stability suggested that I should conduct the comparison between Mark and Sign by reference only to Signs appearing in Annex 8I in relation to the relevant Model in issue and it produced in its skeleton the only other six Getty Images watermarks* included in Annex 8I:



417. I bear in mind Getty Images' criticism that these were artificially 'blown up' and presented out of context, but, to my mind, what is clear is that none of these alternative Signs is identical to the Marks and I do not consider that they would be regarded as identical by an average consumer. The top left watermark* was affixed to the Alien Landscape Image to which I have already referred and is so blurred that even Getty Images conceded in closing that it cannot be relied upon. The remaining five all include obvious additional letters, or (in one case) an 'm' with three arches. On balance, none of these things is so insignificant that it would go unnoticed by the average consumer. A couple would come within the category of "nonsensical" watermarks* referred to by Mr Auerhahn. Although Getty Images stated in opening that, in the event that Stability sought to rely on any other examples it would address them in closing, I did not understand it to make any observations about any of these additional examples.

418. I agree with Stability that, in the circumstances, Getty Images has not relied upon a single example of a representative Sign (bearing a Getty Images watermark*) which is identical to the Marks in issue for *any* of the Models in respect of which I have found that a watermark* will have been generated in real life by a user in the UK.

419. In closing, Ms Bowhill submitted that if the court was going to look at watermarked images going outside the four images ordered by the court, it was only fair that it should do so having regard (also) to all of the Signs in Annex 8H (which runs to 101 pages). However, I do not consider this to be appropriate. I have held that there is no evidence of a real world user making use of verbatim, re-worded or "inspired" prompts and so it would be wrong to look at these.

420. There is no need to look at other Signs generated using the “vector art” prompt as I have found identity having regard to the examples chosen by Getty Images.

421. That leaves only images generated using the words “news photo”. Getty Images made no attempt at trial to identify any watermarks* on these images on which they specifically sought to rely for this purpose and I cannot see that it is for the court to proceed to carry out the necessary comparison without appropriate submissions from each side. This is particularly so where, as I have said above, the enquiry as to use is highly fact sensitive and depends upon an analysis of both the synthetic image and the watermark* appearing on that image. If Getty Images had wished to draw my attention specifically to any of the images in Annex 8H as an alternative to their chosen images, they could have done so in closing. Following circulation of the draft judgment, Getty Images suggested that the failure in the judgment to consider all other “news photo” images bearing the Getty Images watermark* in Annex 8H was a material omission to which they had a duty to draw my attention. They invited me to consider 17 additional images in Annex 8H in connection with their claims under sections 10(1) and 10(2) TMA. Alternatively they invited me to permit both parties to make further submissions on the point. Having considered the matter with care, I decline to take this course. Given that Getty Images had every opportunity at trial to make submissions about any of the images in Annex 8H, I cannot see that it is an obvious omission that this judgment does not deal with images in respect of which they chose to make no submissions whatsoever (with the obvious consequence that Stability also made no submissions on those images). I do not consider a very general submission made in closing to the effect that I should consider images in Annex 8H (as recorded in paragraph 419 above) to be sufficient. It is far too late to invite the court after circulation of the judgment in draft to consider evidence in respect of which no submissions were made at trial. In a trade mark case involving comparison between mark and sign (and particularly in this case where all signs are rendered differently and appear on synthetic images of many different types), the court will expect to receive proper submissions on any relevant sign upon which reliance is placed for the purposes of the comparison. It is unsatisfactory and inappropriate to expect the court to undertake such an exercise in the absence of any submissions whatever (including by way of identification of any specific images on which reliance is placed). That a party may appreciate too late that it should have focused its attention elsewhere in the evidence is not a reason to consider that an omission has been made in the judgment and nor is it a reason to permit a party to re-open issues that have been fully argued at trial. It has been said before that the trial is not a dress rehearsal and that all and any arguments that a party wishes to make should be made at the trial. Getty Images had every opportunity to make out their case at trial and it would not now be consistent with the overriding objective to permit them to re-open that case.

Conclusion on identity of Mark and Sign

422. I have found identity of Mark and Sign in relation to the ISTOCK Marks. There is, however, no evidence of any real world use of a sign which is identical to the Getty Images Marks.

Identity of Goods and Services

423. Getty Images pleads that Stability has used the Signs in relation to “synthetic image outputs” and “the provision of synthetic image outputs” which are identical goods and

services to goods and services specified in classes identified in each of the individual Marks. I consider this topic for both the Getty Images Marks and the ISTOCK Marks, notwithstanding that, if I am right to dismiss the identity of Mark and Sign in relation to the Getty Images watermark* examples, then this analysis is strictly unnecessary in respect of the use of that sign in those examples for the purposes of the section 10(1) TMA claim.

424. Although Stability accepts similarity between its goods and services and those registered for the Marks, it denies that they are identical. This is primarily because it contends that none of the goods and services for which the Getty Images Marks are registered can be construed as at the date of registration in such a way as to be identical to Stability's goods and services – this, it says would be an impermissible exercise in hindsight. It submits that although the language of the specification may be very general, its scope cannot be construed to cover additional activities unforeseen at the time it was filed.
425. I begin by considering some general points raised by the parties.
426. Stability accepts that the words “synthetic image outputs” are an appropriate characterisation of the intangible outputs of the Stable Diffusion Model. In closing, Getty Images appeared to suggest that Stability’s goods might in fact be characterised as “photographs” on the basis that that is how some users would refer to the outputs of the Model. I reject that case. It is not pleaded and it is far too late to advance a new characterisation of the goods in question only in closing. A synthetically generated image is not a photograph and I do not accept that it would be regarded as such for the purposes of trade. The Model is certainly not taking or composing photographs in the literal meaning of those words. That seems to me to be so even if the words ‘photo’ or ‘photograph’ might sometimes be used (wrongly) to describe such an image (see by analogy *British Sugar* – jam is not a dessert sauce even though “it too can be used on a dessert and everyone knows and sometimes does this” (at page 289, per Jacob J)).
427. Stability contends that the phrase “provision of synthetic image outputs” is “legalese” and does not reflect real world usage. It submits that such language fails to recognise that the user of the service is guiding the generation of images and controlling the contents of the outputs using their individual prompts. It says that a distinction must be drawn between the various access mechanisms: when the Model is downloaded from GitHub and Hugging Face, the user gets an AI image generator in the form of source code and model weights which are intangible goods, whereas, by contrast, DreamStudio and the Developer Platform are Generative AI services.
428. The question of whether the code and model weights are “intangible goods” is a question to which I shall need to return in the context of the case on secondary infringement of copyright. I do not understand it to be a question that I need to determine in connection with this issue, and, in any event, the code and model weights are different “goods” from the “synthetic image outputs” pleaded by Getty Images and accepted by Stability as a fair characterisation of the Stability outputs. Getty Images does not plead or pursue a case that the code and model weights made available to users for the purposes of download to a local computer is an identical (or similar) “service”.
429. Getty Images does not distinguish between the different access mechanisms, essentially because its case is that for each mechanism there is both the provision of the identical

service and the provision of identical goods. I accept for DreamStudio and the Developer Platform (for reasons explained in more detail below) there is both. However, I agree with Stability that the real nature of the service in relation to the ‘download’ mechanism is releasing the code and model weights to users. I therefore agree with Stability that there is no identity between this service and the relevant specifications upon which Getty Images relies. However, to my mind, the goods then generated by the Model are synthetic image outputs – and this applies to all access mechanisms.

430. I reject Stability’s case that the word ‘provision’ is legalese and does not reflect real world usage. I have already made findings as to Stability’s level of control and these findings are also relevant here. By releasing the Models, Stability has made them (or more accurately the software which makes up the Model) available to the public to use – specifically, for this purpose, via the DreamStudio platform and via an API in the form of the Developer Platform. It is clear from the Expert evidence that the Model does not itself store images and so does not “provide” the user with an image that already exists – instead it generates a new image. But, although it is obviously more accurate to say that the Model *generates* synthetic images, I do not consider it to be unrealistic or unrepresentative of real world use to infer that many consumers will regard the generation of synthetic images by DreamStudio or the Developer Platform as involving “the provision” of such outputs. That seems to me to be particularly the case where, as Stability accepts, DreamStudio and the Developer Platform are themselves services. Furthermore, as Ms Bowhill points out (and I agree), the act of creation is not an act of trade mark infringement and so this is something of an artificial point in any event. The court must focus on the acts which constitute trade mark use – here the affixing of the Sign to an image.

431. Getty Images plead that these goods (synthetic image outputs) and services (the provision of synthetic image outputs) are identical to various of the goods and services for which the Marks are registered. Their best case on this was first identified in a letter dated 28 May 2025 and then set out in tables in their opening submissions dealing with the Getty Images Marks and the ISTOCK Marks as follows:

Trade Mark	Goods/Services
UK859: GETTY IMAGES	Class 9: digital media, digital materials, digital content, ... namely images...photographs...news images...and the above mentioned goods stored or recorded on electronic or computer media or downloadable from databases or other facilities provided over global computer networks, wide area networks, local area networks, or wireless networks Class 42: computerized on-line search and retrieval services for

	images...photographs...still images...news images
UK005: GETTY IMAGES	<p>Class 9: computer software on ... downloadable format for use in the fields of creating and manipulating visual media, graphic images, news images...</p> <p>Class 16: photographs</p> <p>Class 35: retail sales of photographs, namely stock photography services</p> <p>Class 35: online retail services in the field of art, namely...photographs</p> <p>Class 38: electronic delivery of images, photographs...graphic images...news images...via a global computer network and other computer networks</p> <p>Class 38: providing access to various media, namely digital stock photography archival photographs...news images...via an interactive computer database</p> <p>Class 38: computer aided transmission of ...images</p> <p>Class 41: digital imaging services</p> <p>Class 41: photographic computer imaging</p> <p>Class 42: providing access to an interactive online computer database in the fields of visual media, graphic images...photography...</p> <p>Class 42: providing access to various media namely digital stock photography, archival photography...news images... via an interactive computer database</p>
UK925: gettyimages	<p>Class 16: photographs</p> <p>Class 41: digital imaging services</p>

Trade Mark	Goods/Services
UK297: ISTOCK	<p>Class 40: photocomposing services</p> <p>Class 41: digital imaging services</p>

UK819: ISTOCK	Class 9: downloadable digital illustrations and graphics
---------------	--

432. Without limiting her clients' case, Ms Lane submitted in closing that synthetic image outputs and/or the provision of synthetic image outputs plainly fall within:

“(i) digital media, digital materials, digital content, ... namely images...and the above mentioned goods stored or recorded on electronic or computer media or downloadable from databases or other facilities provided over global computer networks” (class 9); (ii) “electronic delivery of images...via a global computer network and other computer networks” (class 38); (iii) “digital imaging services” (class 41); and (iv) “downloadable digital illustrations and graphics” (class 9).

She did not attempt to justify any of the other pleaded specifications.

433. I am bound to say that Stability's oral submissions on this issue in closing were a great deal more extensive than its written submissions. Many points were raised for the first time without notice to Getty Images and often appeared to be straying into the territory of questioning the validity of the specifications – not an issue in these proceedings. If Stability intended to take wide-ranging points about the construction of the specification at trial, I consider that it should have alerted Getty Images to those points much earlier than in its oral closing submissions. Furthermore, in so far as Stability sought to distinguish in those submissions between various access mechanisms, it should also have made that clear much earlier.

434. Notwithstanding Mr Cronan's engaging submissions, I have not always found it easy to follow how it is said that the specifications apply differently to different access mechanisms. My task was not made any easier by the fact that Getty Images (no doubt because they had not understood the extent of the dispute on this topic) dealt with the issue of double identity with what might be described as an extremely light touch, often not seeking to explain what individual specifications meant, but instead relying on little more than bald assertions that those specifications covered the goods and services in respect of which Stability had used the signs.

435. Owing to Getty Images' continuing case that I must look at all of the goods and services relied upon against each Mark (and not simply the goods and services on which they focused in closing), I deal with them all below. In so doing, I also address Stability's arguments, as I understood them, as best I can.

UK297:

436. I do not consider the provision of synthetic image outputs to be identical to “photocomposing services” (class 40). The entire extract relied upon in the Mark from Class 40 reads “Treatment of materials; Bookbinding; framing of works of art; lithographic printing; photocomposing services; photographic printing; Printing”. Stability provided me in closing with an extract from the Encyclopaedia Britannica identifying the meaning of photocomposition as a “method of assembling or setting type by photographing characters on film from which printing plates are made. The

characters are developed as photographic positives on film or light sensitive paper from a negative master containing all the characters; the film, carrying the completed text, is then used for making a plate for gravure or lithographic printing by a photomechanical process". Bearing in mind this definition (which I did not understand to be disputed by Getty Images) and its context in the specification, it seems to me that Getty Images was right not to rely specifically upon this in its oral closing submissions.

437. However, on balance I consider the provision of synthetic image outputs (on DreamStudio and via the API) to fall within the words "digital imaging services" (class 41) and to be consistent with their core meaning. In closing Getty Images made no attempt to explain what this term might mean or why it was identical in the case of one or more of the access mechanisms. Stability also produced nothing designed to assist on the meaning of the words in the specification. Mr Cronan submitted that "digital imaging services" are vague words which, as at 2003 (i.e. the date of registration of the Mark) would have meant editing and processing of digital images – although he produced nothing in support of this submission.

438. Doing the best I can, it seems to me that there is a lack of precision in the term "digital imaging services" at least in the sense that it is not clear what such services might cover. Attempting to confine it to the core or substance of its possible meaning (see *SkyKick* at [365]) appears to me to involve construing it to mean the creation of a digital representation of the visual characteristics of an object, as well as the editing and processing services suggested by Stability. No doubt the development of Stable Diffusion provides a new means of achieving this, but the core service appears to me to be the same: the use of software online (DreamStudio) or via an API to facilitate the creation of digital images. Indeed this appears to me to be an example of "a new variant" of a service which falls within the specification, of the type envisaged in *Reed v Reed*. It is not a different service (cf. the decision in *SkyKick* to the effect that Cloud Migration is not an electronic mail service). I therefore consider there to be identity in relation to this phrase in so far as access via DreamStudio and the API is concerned. Downloading the Model via GitHub and Hugging Face appears to me to be somewhat different because users are able to generate the images locally and the service is in fact the provision of access to the Model in the first place, thereby enabling the user to generate images locally.

UK819:

439. I consider that synthetic image outputs of the "vector art" type fall within the words "downloadable digital illustrations and graphics" (class 9). Even Stability accepts (in its written submissions) that these words are clear and that such "vector art" type 2D image outputs "are very close to being referable as digital illustrations", albeit that it suggests that "on balance, it would not be normal use of language in relation to these models to say "I generated a digital illustration"". In closing, Stability sought to run an additional argument to the effect that a full extract from the specification reads "...downloadable digital photographs, illustrations, audio clips, video clips, fonts, code snippets and graphics...", submitting that none of digital "photographs", "illustrations" or "graphics" is identical to a synthetic image output. On balance I disagree. I do not consider the specification to include ambiguous terms and I consider that the average consumer might well say that he or she had generated a digital illustration when describing an output image generated by Stable Diffusion of the type to which the

iStock sign has been affixed. I therefore consider there to be identity between this specification and use of the Signs generated by each of the three access mechanisms.

440. Turning to the Getty Images Marks:

UK859:

441. I take first “digital media, digital materials, digital content, … namely images… and the above mentioned goods stored or recorded on electronic or computer media or downloadable from databases or other facilities provided over global computer networks…” (class 9). Once again, Getty Images advance no clear description of what this is to be taken to mean, asserting in their closing submissions merely that synthetic output images fall “very easily” into this class. Stability submits that “digital media” and “images” are vague terms, which, pared back to their core meaning as at the date of registration, can only include stock photographs, editorial photographs and illustrations. I tend to disagree. I do not consider “digital media” or “images” to be vague or ambiguous and it seems to me that these terms plainly cover synthetic image outputs.

442. Stability also seeks to construe the words “other facilities” by reference to the word “database”, contending that “other facilities” must therefore mean “some kind of pre-arranged collection of content on-line” and that this must really be describing “how a stock content website works in 2012” when the Mark was registered. Again I disagree. While an AI model is not a “database”, it has been trained from a database and is another “facility” provided over a global computer network.

443. Finally, Stability seeks to differentiate between the different access mechanisms by reference to the concepts of “storage” and “downloading”. There is nothing in these points. In DreamStudio, files are stored on a server after generation (as Stability accepts) and they are also downloadable “from databases or other facilities provided over global computer networks”. When using the API, images are provided to the user via an input/output service (effectively, as I understand it, a means of making images available to the user without requiring a “download”) and when using the Model on a local computer (having downloaded it from GitHub and Hugging Face), images are generated on the local computer where they may be “stored or recorded”. Synthetic output images generated by each of the three access mechanisms thus appear to me to be new variants of the goods described in the specification (*Reed v Reed*).

444. I reject the proposition that the provision of synthetic image outputs falls within “computerized on-line search and retrieval services for images … photographs … still images…news images” (class 42). The provision of new images using text prompts does not appear to me to be an aspect or feature of a search or retrieval service for existing images (see *SkyKick* at [366]).

UK005:

445. Getty Images say that the provision of synthetic image outputs also falls within “electronic delivery of images…via a global computer network or other computer networks” (class 38). Class 38 is concerned mainly with services that allow communication together with the broadcasting and “transmission of data”. The Explanatory Note to Class 38 states that the class does not include “content or subject

matter that may be contained in the communication activity...”. I agree with Stability that this class therefore cannot be identical to synthetic image outputs, because it is not concerned with goods, but I consider that it is plainly capable of encompassing the provision of synthetic image outputs using the DreamStudio platform and the Developer Platform – both of which involve the electronic delivery of images via a computer network. To my mind, this specification cannot, however, encompass the ‘download’ access mechanism because, when the Model is being used on a local computer, the synthetic output images are being generated locally – they are not being delivered via a global network or any other computer network.

446. For reasons set out above, I consider that the provision of synthetic image outputs are identical to “digital imaging services” (class 41) when they are being generated by the user on the DreamStudio platform or via the Developer Platform.
447. As identified in the table above, Getty Images pleads a number of other specifications in relation to this Mark, under classes 9, 16, 35, 38, 41 and 42. Dealing with these as concisely as possible:
 - i) the class 9 specification is concerned with “computer software” as the relevant “good” which is not pleaded and is not identical to synthetic output images;
 - ii) class 16 “photographs” are not synthetic image outputs for reasons I have explained;
 - iii) the class 35 specification is concerned with “retail sales of photographs, namely stock photography services”, which is plainly not identical to the provision of synthetic image outputs; and “online retail services in the field of art, namely...photographs” again not identical to the provision of synthetic image outputs. An AI model does not sell photographs online.
 - iv) The class 38 specification is concerned with “providing access to various media, namely digital stock photography archival photographs...news images...via an interactive computer database” and “computer aided transmission of ...images”. The former (even assuming synthetic images to fall within the types of images identified), presupposes the provision of access via a “computer database”, something which Getty Images expressly plead that the model weights are not. The latter appears to assume the transmission of existing images by computer and (aside from the fact that this is obviously imprecise and it is unclear what core service this might relate to) the provision of synthetic image outputs is not identical.
 - v) Class 41: photographic computer imaging is not the provision of synthetic output images.
 - i) The class 42 specifications relied upon: “Class 42: providing access to an interactive online computer database in the fields of visual media, graphic images...photography...” and “providing access to various media, namely digital stock photography, archival photography...news images... via an interactive computer database”, both rely upon the existence of an interactive database, which an AI model is not. Furthermore, downloading the Model via GitHub and Hugging Face takes the model offline so that there is no access to images being provided “online”.

448. Getty Images made no attempt in closing to explain how any of the specifications referred to above might be identical and, accordingly, I reject its case on identity in respect of these specifications.

UK925

449. Again there appears to me to be identity between “digital imaging services” (class 41) and the provision of synthetic output images by the DreamStudio platform and via the Developer Platform for the reasons set out above.

450. I reject Getty Images’ case of identity in respect of “photographs” (class 16), again for reasons I have given.

Conclusion on identity of goods and services

451. For the reasons set out above there is:

- i) identity of goods (but not services) in relation to the ‘download’ access mechanism (a) in respect of the Getty Images Marks (Models v1.x and v2.x); and (b) in respect of the ISTOCK Marks (Model v1.x only); and
- ii) identity of goods and services in relation to the use of DreamStudio and the Developer Platform in respect of the Getty Images Marks (Models v1.x and 2.x); and (b) in respect of the ISTOCK Marks (Model v1.x only).

452. As appears from the above analysis. Getty Images continued to rely at trial on a number of specifications for various of the Marks in respect of which they advanced no positive case. I consider these should properly have been abandoned. It is neither desirable nor consistent with the overriding objective to require the court to consider parts of a specification absent any properly articulated case.

Conclusion on section 10(1) Infringement

453. For the reasons set out above:

- i) I find double identity infringement by Stability in respect of iStock watermarks* generated by users of v1.x (accessing v1.x via the API and accessing v1.4 through DreamStudio). This finding is based specifically on the example watermark* shown on the Spaceships Image. There is no evidence as to the Model that generated the Dreaming Image. It is impossible to know how many (or even on what scale) watermarks* have been generated in real life that would fall into a similar category.
- ii) I dismiss the claim of double identity infringement in relation to v1.x and v2.x in respect of the Getty Images watermarks*.
- iii) I dismiss the claim of double identity infringement in relation to v2.x in respect of the iStock watermarks*.
- iv) In circumstances where I have determined that there is no evidence of a user in the real world generating an image bearing a watermark from either SD XL or v1.6, I dismiss the claim of double identity infringement in relation to those

Models in respect of both the Getty Images watermarks* and the ISTOCK watermarks*.

SECTION 10(2) INFRINGEMENT

The Law

454. Section 10(2)(b) provides, in so far as relevant that:

“(1) ...

(2) A person infringes a registered trade mark if he uses in the course of trade a sign where because—

(a) ...

(b) the sign is similar to the trade mark and is used in relation to goods or services identical with or similar to those for which the trade mark is registered,

there exists a likelihood of confusion on the part of the public, which includes the likelihood of association with the trade mark.”

455. In order to establish infringement under s 10(2), a claimant must satisfy six requirements: (i) there must be use of a sign by a third party within the relevant territory; (ii) the use must be in the course of trade; (iii) it must be without the consent of the proprietor; (iv) it must be of a sign which is at least similar to the trade mark; (v) it must be in relation to goods or services which are at least similar to those for which the trade mark is registered; and (vi) it must give rise to a likelihood of confusion: *Muzmatch* per Arnold LJ at [26].

456. Conditions (i)-(iii) are dealt with above.

457. Stability has conceded similarity for the purposes of (v) on Getty Images’ ‘best case’ goods and services – i.e. those expressly relied upon in closing submissions. Stability also does not contest that Getty Images can overcome the threshold requirement of similarity between Mark and Sign (condition (iv)). Accordingly an assessment of the degree of similarity becomes relevant to the question of whether “there exists a likelihood of confusion on the part of the public” (*Iconix* at [32]).

458. Given my findings above, Stability’s concession of similarity applies to both the iStock watermark* images and to the Getty Images watermark* on the First Japanese Temple Garden Image. The watermark* on the Obama Image is not representative and so I address it no further, other than to say that if I am wrong about that, then it is plainly highly similar.

459. I find that, even if I am wrong as to the iStock Signs on the Spaceships Image and the Dreaming Image being identical, they are nevertheless highly similar to the ISTOCK Marks. Although I have found that the Sign on the First Japanese Temple Garden Image is not identical, I find for the purposes of section 10(2) TMA that it is highly similar to the Getty Images Marks. That seems to me to be the impression that would be created

on the mind of the average consumer (always bearing in mind that the average consumer rarely has the chance to make direct comparisons between mark and sign). I did not understand Stability to make any alternative submission.

Likelihood of confusion: the global assessment

460. The Trade Marks Registry has adopted a standard summary of the principles established by these authorities for use in the registration context, which has been approved by the Court of Appeal on a number of occasions. The most recent version of this summary was given by Arnold LJ in *Muzmatch* at [27], and approved by the Supreme Court in *Iconix* at [38]):

- “(a) the likelihood of confusion must be appreciated globally, taking account of all relevant factors;
- (b) the matter must be judged through the eyes of the average consumer of the goods or services in question, who is deemed to be reasonably well informed and reasonably circumspect and observant, but who rarely has the chance to make direct comparisons between marks and must instead rely upon the imperfect picture of them he has kept in his mind, and whose attention varies according to the category of goods or services in question;
- (c) the average consumer normally perceives a mark as a whole and does not proceed to analyse its various details;
- (d) the visual, aural and conceptual similarities of the marks must normally be assessed by reference to the overall impressions created by the marks bearing in mind their distinctive and dominant components, but it is only when all other components of a complex mark are negligible that it is permissible to make the comparison solely on the basis of the dominant elements;
- (e) nevertheless, the overall impression conveyed to the public by a composite trade mark may, in certain circumstances, be dominated by one or more of its components;
- (f) and beyond the usual case, where the overall impression created by a mark depends heavily on the dominant features of the mark, it is quite possible that in a particular case an element corresponding to an earlier trade mark may retain an independent distinctive role in a composite mark, without necessarily constituting a dominant element of that mark;
- (g) a lesser degree of similarity between the goods or services may be offset by a greater degree of similarity between the marks, and vice versa;

- (h) there is a greater likelihood of confusion where the earlier mark has a highly distinctive character, either *per se* or because of the use that has been made of it;
- (i) mere association, in the strict sense that the later mark brings the earlier mark to mind, is not sufficient;
- (j) the reputation of a mark does not give grounds for presuming a likelihood of confusion simply because of a likelihood of association in the strict sense; and
- (k) if the association between the marks creates a risk that the public might believe that the respective goods or services come from the same or economically-linked undertakings, there is a likelihood of confusion.”

461. As I have already alluded to, it is well established, and has recently been confirmed by the Supreme Court in *Iconix*, that it can be relevant to take into account the post-sale context when considering trade mark issues and that post sale confusion amounts to a likelihood of confusion. The rationale is that a mark may still function as a trade mark after sale because it operates as a badge of origin, and hence quality, after the goods are sold. It does so not primarily to the purchaser of the goods (who is likely to be aware of their origin), but to third parties who encounter the goods after sale (see *Montres Breguet* (CA) at [84]-[85] and *Iconix* at [75]).
462. In this case the Signs will be encountered after the user has downloaded the Model, or has used one of the other access mechanisms described in Appendix B. However, this does not preclude the scope for confusion. Those users will be encountering the Sign in a realistic and representative way when it appears on their screens. In some cases, the user may encounter the Sign shortly after he or she has used one of the access mechanisms. In other cases, he or she may encounter the Sign long after using one of the access mechanisms.
463. I have already addressed the question of the available evidence in this case. In the context of considering confusion, I bear in mind that the absence of evidence is hardly ever determinative and that claims under section 10(2) TMA can succeed in the absence of evidence.
464. In my judgment, the global assessment points to a likelihood of confusion on the part of a significant proportion of the relevant public as a result of the appearance of watermarks* (in the form of those appearing on the First Japanese Temple Garden Image, the Spaceships Image and the Dreaming Image) for the following reasons.
 - i) It is common ground that the Marks have a high degree of distinctive character.
 - ii) There is a very high degree of similarity between Mark and Sign having regard to these specific examples.
 - iii) There is identity (or, if not identity, a high degree of similarity) between goods and services.

- iv) Two classes of average consumer (downloading the Model and accessing it via the Developer Platform) will be technically savvy and will be paying at least a moderate degree of attention. Owing to their understanding of how AI models work, including their understanding of how they are trained, they are unlikely to believe that Getty Images is supplying the synthetic images but, as I have found, a significant proportion will think that there is a connection or material link between Stability and Getty Images, including because they will think that Stable Diffusion was trained on Getty Images Content “under license” from Getty Images. The October 2022 SalesForce Message to which I have referred above provides support for such confusion in the real world.
- v) Average consumers accessing the Model using DreamStudio will be less technically astute. A significant proportion, paying a moderate degree of attention, will think that a generated image bearing a watermark* has been supplied by Getty Images or that Stable Diffusion has been trained on Getty Images Content under license from Getty Images, or that there is some other economic link between Getty Images and Stability. I agree with Getty Images that the average consumer in this class (or a significant proportion) would not assume that a major corporation such as Getty Images would allow its assets and signs to be used by third parties without its permission. Even assuming that such consumer was not confused as to the origin of the image, the natural assumption would be that the synthetic output was generated by a company that had some form of licensing or other economic arrangement in place with Getty Images – there would thus be confusion as to the licensing position and thus the extent of any association with the Marks.
- vi) Whatever the access mechanism, I do not consider the context in which the average consumer encounters the Signs to undermine this conclusion for all the reasons I have already explained, albeit in a slightly different context - Stability accepted in closing that the relevant factors going to “use in relation to goods and services” are the same as those relied upon in connection with the likelihood of confusion. The average consumer will not think that he or she has full control over the generation of images bearing watermarks* or that he or she bears responsibility for the appearance of the watermarks*. Although the levels of technical sophistication amongst the classes of average consumer will be different, I find a likelihood of confusion in relation to each class (or a significant proportion of each class) for all the reasons I have given.
- vii) Furthermore, I do not consider that Stability’s argument as to the “garbled nature of the sign” assists for reasons set out above. The global assessment is highly fact sensitive and the specific Signs with which I am concerned are not so garbled as to alert the average consumer to an issue.
- viii) As I have found, many consumers will be familiar with “the regime” of using watermarks* (as described by Ms Cameron) to indicate that Content must be paid for, but this does not preclude the likelihood of confusion over the appearance of images displaying a Getty Images or an iStock watermark*.
- ix) It does not matter that the Sign would only be encountered in a post-sale context. That context is both realistic and representative.

Conclusion on section 10(2) Infringement

465. In light of these findings, I consider Getty Images to have made out its case of section 10(2) TMA infringement in relation to:

- i) Model v1.x in respect of the ISTOCK Mark (for users accessing (i) v1.x via the Developers Platform and (ii) v1.4 through DreamStudio).
- ii) Model v2.1 in respect of the Getty Images Mark.

466. However, the above global assessment applies specifically to the three examples of use of the Sign which I have found to be identical or similar (the Spaceships Image, the Dreaming Image and the First Japanese Temple Garden Image – albeit that it is unclear which Model generated the Dreaming Image). For reasons I have explained, the analysis is highly fact sensitive – while Stability’s points on the garbled nature of the watermarks* and the nature of the images produced do not appear to me to weigh heavily in the balance on the global assessment of these three examples, it might be the case that they would do so in respect of different watermarks* generated on different images in response to different prompts. That this must be the case is borne out by a consideration of various of the examples in both Annex 8H and Annex 8I. I should add that I have not forgotten that Getty Images watermarks* have been produced by v1.x using the “news photo” prompt, but absent specific submissions directed at those watermarks* and the images on which they appear, I do not see that I can properly make a finding of section 10(2) infringement in relation to them, and (until after circulation of the draft judgment) I was not invited to do so. For reasons I have already explained, I do not consider it to be consistent with the overriding objective to permit Getty Images to have a second bite of the cherry once the judgment has been circulated.

467. Unlike the more usual case of trade mark infringement, it is impossible to conclude that for every watermark* generated by the same version of the Model, a similar analysis will apply such that infringement would follow automatically. The question of how many similar (infringing) watermarks* are likely to have been generated by users in the UK remains unknowable and Getty Images has advanced no case at this trial designed to address that question.

SECTION 10(3) INFRINGEMENT

The Law

468. Section 10(3) TMA provides as follows:

“(3) A person infringes a registered trade mark if he uses in the course of trade, in relation to goods or services, a sign which—

- (a) is identical with or similar to the trade mark,
- (b)

where the trade mark has a reputation in the United Kingdom and the use of the sign, being without due cause, takes unfair advantage of, or is detrimental to, the distinctive character or the repute of the trade mark.”

469. Section 10(3A) provides that “[s]ubsection (3) applies irrespective of whether the goods and services in relation to which the sign is used are identical with, similar to or not similar to those for which the trade mark is registered”.

470. The aim of section 10(3) is to protect the reputation of a registered trade mark from damage in the form of unfair advantage or detriment to the distinctive character or repute of the mark. There is no requirement for confusion and no requirement for the infringing use to be in relation to identical or similar goods or services.

471. There are nine conditions that a claimant must satisfy in order to establish infringement under section 10(3): (i) the trade mark must have a reputation in the relevant territory; (ii) there must be use of a sign by a third party within the relevant territory; (iii) the use must be in the course of trade; (iv) it must be without the consent of the proprietor of the trade mark; (v) it must be of a sign which is at least similar to the trade mark; (vi) it must be in relation to goods or services; (vii) it must give rise to a “link” between the sign and the trade mark in the mind of the average consumer; (viii) it must give rise to one of three types of injury, that is to say, (a) detriment to the distinctive character of the trade mark, (b) detriment to the repute of the trade mark or (c) unfair advantage being taken of the distinctive character or repute of the trade mark; and (ix) it must be without due cause: See *Muzmatch* per Arnold LJ at [55].

472. In this case, condition (i) is admitted and conditions (ii) to (vi) are dealt with above. Condition (vii), link, is not contested. No point is taken on condition (ix). Accordingly, the only remaining issue is whether the use of the Sign gives rise to one of the three types of injury (condition (viii)), namely, detriment to distinctive character (dilution), detriment to reputation (tarnishment) and unfair advantage (free-riding), all of which are pleaded by Getty Images.

473. In opening submissions, there appeared to be a dispute between the parties as to (i) exactly what is required in order to establish one or more of the three types of injury (specifically whether a change in economic behaviour is required for each); and (ii) whether Getty Images’ pleaded case in relation to each of the different types of injury is sufficient.

474. Stability relies upon an observation made by Arnold LJ in *Thatchers Cider Company Ltd v Aldi Stores Ltd* [2025] EWCA Civ 5 (“*Thatchers*”) at [51] to the effect that:

“the case law of the Court of Justice establishes that infringement under section 10(3) requires evidence of a change in the economic behaviour of the average consumer of the goods or services for which the trade mark is registered”.

475. Stability says that this observation (made in the context of a case concerned with unfair advantage and detriment to reputation) applies to each of the types of injury under s.10(3) and that Getty Images have pleaded change in economic behaviour in relation only to an allegation of detriment to the distinctive character of the Marks. It submits that this single allegation of change in economic behaviour cannot succeed (for reasons to which I shall return). Further it says that Getty Images cannot succeed on either of the other two heads of injury asserted, not least because there is no plea of change in economic behaviour in relation to each of those heads.

476. By the time of closing submissions, I did not understand Getty Images seriously to suggest that Arnold LJ's observation in *Thatchers* could properly be limited only to dilution. Instead, it made a rather different point, to the effect that Arnold LJ's observation did not address the more nuanced approach to evidential requirements taken by the CJEU (to the effect that a deduction may be made based on "an analysis of the probabilities and by taking account of the normal practice in the relevant commercial sector as well as the other circumstances of the case": Case C-383/12 *Environmental Manufacturing LLP v Office for Harmonisation in the Internal Market (Trade Marks and Designs)* [EU:C:2013:741]).

477. As for its pleaded case, Getty Images accept that their Particulars of Claim make a plea of change in economic behaviour only when dealing with dilution of the distinctiveness of the Marks. However they contend that there is no requirement for an express plea in relation to the allegations of free-riding and tarnishment because "[t]he serious risk of any detriment and the relevant change in the economic behaviour of consumers can be deduced from the facts and matters that are pleaded in those sub-paragraphs".

478. I begin by taking the applicable legal principles on each type of injury in turn.

Detriment to distinctive character/Dilution

479. In *Lidl* at [22]-[24], Arnold LJ set out the following propositions in relation to detriment to distinctive character with reference to C-252/07 *Intel Corp Inc v CPM UK Ltd* EU:C:2008:655:

"The Court of Justice stated in *Intel v CPM*:

"29. As regards, in particular, detriment to the distinctive character of the earlier mark, also referred to as 'dilution', 'whittling away' or 'blurring', such detriment is caused when that mark's ability to identify the goods or services for which it is registered and used as coming from the proprietor of that mark is weakened, since use of the later mark leads to dispersion of the identity and hold upon the public mind of the earlier mark. That is notably the case when the earlier mark, which used to arouse immediate association with the goods and services for which it is registered, is no longer capable of doing so.

...

67. The more immediately and strongly the earlier mark is brought to mind by the later mark, the greater the likelihood that the current or future use of the later mark is taking unfair advantage of, or is detrimental to, the distinctive character or the repute of the earlier mark.

68. It follows that, like the existence of a link between the conflicting marks, the existence of one of the types of injury referred to in Article 4(4)(a) of the Directive, or a serious likelihood that such an injury will occur in the future, must be

assessed globally, taking into account all factors relevant to the circumstances of the case, which include the criteria listed in paragraph 42 of this judgment.

69. ... the stronger the earlier mark's distinctive character and reputation the easier it will be to accept that detriment has been caused to it ...

...

77. ... proof that the use of the later mark is or would be detrimental to the distinctive character of the earlier mark requires evidence of a change in the economic behaviour of the average consumer of the goods or services for which the earlier mark was registered consequent on the use of the later mark, or a serious likelihood that such a change will occur in the future.

78. It is immaterial, however, for the purposes of assessing whether the use of the later mark is or would be detrimental to the distinctive character of the earlier mark, whether or not the proprietor of the later mark draws real commercial benefit from the distinctive character of the earlier mark.”

23. With respect to the requirement identified in *Intel v CPM* at [77], the Court of Justice added in Case C-383/12 *Environmental Manufacturing LLP v Office for Harmonisation in the Internal Market (Trade Marks and Designs)* [EU:C:2013:741]:

“42. Admittedly, Regulation No 207/2009 and the Court’s case-law do not require evidence to be adduced of actual detriment, but also admit the serious risk of such detriment, allowing the use of logical deductions.

43. None the less, such deductions must not be the result of mere suppositions but ... must be founded on ‘an analysis of the probabilities and by taking account of the normal practice in the relevant commercial sector as well as all the other circumstances of the case’.”

24. It is not in dispute that the approach articulated in [43] is also applicable to the question of whether there has already been a change to the economic behaviour of the average consumer.

480. It is plain from this passage that, when Arnold LJ was referring to the need for evidence in *Thatchers*, he was using a shorthand. The existence of a change in the economic behaviour of the average consumer, or a serious likelihood of such a change, may be deduced having regard to the factors identified. I did not understand Stability to disagree with this proposition.

Detriment to Repute

481. In *Thatchers* at [50] the Court of Appeal cited with approval the definition given by the Court of Justice in C-487/07 *L'Oréal SA v Bellure NV* EU:C:2009:378 (“*L'Oréal v Bellure*”) at [40]:

“As regards detriment to the repute of the mark, also referred to as ‘tarnishment’ or ‘degradation’, such detriment is caused when the goods or services for which the identical or similar sign is used by the third party may be perceived by the public in such a way that the trade mark’s power of attraction is reduced. The likelihood of such detriment may arise in particular from the fact that the goods or services offered by the third party possess a characteristic or a quality which is liable to have a negative impact on the image of the mark.”

Unfair Advantage

482. In *Lidl* at [25], the Court of Appeal cited the following paragraphs from *L'Oréal v Bellure*:

“41. As regards the concept of ‘taking unfair advantage of the distinctive character or the repute of the trade mark’, also referred to as ‘parasitism’ or ‘free-riding’, that concept relates not to the detriment caused to the mark but to the advantage taken by the third party as a result of the use of the identical or similar sign. It covers, in particular, cases where, by reason of a transfer of the image of the mark or of the characteristics which it projects to the goods identified by the identical or similar sign, there is clear exploitation on the coat-tails of the mark with a reputation.

...

43. It follows that an advantage taken by a third party of the distinctive character or the repute of the mark may be unfair, even if the use of the identical or similar sign is not detrimental either to the distinctive character or to the repute of the mark or, more generally, to its proprietor.

44. In order to determine whether the use of a sign takes unfair advantage of the distinctive character or the repute of the mark, it is necessary to undertake a global assessment, taking into account all factors relevant to the circumstances of the case, which include the strength of the mark’s reputation and the degree of distinctive character of the mark, the degree of similarity between the marks at issue and the nature and degree of proximity of the goods or services concerned. ... the more immediately and strongly the mark is brought to mind by the sign, the greater the likelihood that the current or future use of

the sign is taking, or will take, unfair advantage of the distinctive character or the repute of the mark ...”.

483. The global assessment may also take into account, where necessary, the fact that there is a likelihood of dilution or tarnishment of the mark (*L'Oréal v Bellure* at [48]).
484. In *Thatchers*, Arnold LJ (having also cited the above passages from *L'Oréal v Bellure*), said this at [48]:

“It is clear both from the wording of the relevant provisions and from the case law of the Court of Justice and General Court interpreting them, in particular *L'Oréal v Bellure*, that this aspect of the legislation is directed at a particular form of unfair competition. It is also clear from the case law both of the Court of Justice and General Court and of the domestic courts and tribunals that the defendant's conduct is most likely to be regarded as unfair where the defendant intends to take advantage of the reputation of the trade mark. Nevertheless, in *Jack Wills Ltd v House of Fraser (Stores) Ltd* [2014] EWHC 110 (Ch), [2014] FSR 39 I concluded at [80] that there was nothing in the case law to preclude the court from holding in an appropriate case that the use of a sign the objective effect of which is to enable the defendant to benefit from the reputation of the trade mark amounts to unfair advantage even if it is not proved that the defendant subjectively intended to exploit that reputation.

485. Paragraph [162] of *Lidl* confirms (in so far as any dispute remains) that a finding of change in economic behaviour, or risk of such change, is required not only for detriment to distinctive character, but also for ‘unfair advantage’.

Getty Images' pleaded case

486. Getty Images' case on injury is set forth in paragraphs 57.7 to 57.11 of their Particulars of Claim, as follows:

“57.7 Such use will result in the Defendant taking advantage of the Claimants' investment in developing the reputation of the Trade Marks and each of them. It is to be inferred that such was the Defendant's intent, alternatively, such is the objective effect of the Defendant's use.

57.8 Further or alternatively, such use will result in the dilution of the distinctiveness of the Trade Marks and each of them. In particular, and pending disclosure and evidence herein, the average consumer of the goods and services for which the Trade Marks are registered are likely to use Stable Diffusion to generate an image which is the same as or similar to the Claimants' Content but in respect of which it will not have to pay a licence fee. Accordingly there is a serious likelihood that the use complained of will result in a change in the economic

behaviour of the average consumer of the goods and services for which the Trade Marks are registered.

57.9 Further or alternatively, Stable Diffusion can be used to create images that contain pornography, violent imagery, and propaganda. Any association with such content will tarnish the reputation of the Trade Marks and each of them. Examples of such images which have been created using Stable Diffusion are at pages 20 and 100 of Annex 8H and Confidential Exhibit DAS-15 to the Witness Statement of David Stanley.

57.10 Further still, the synthetic images generated by Stable Diffusion distort and/or manipulate the underlying image from which it was copied, which is prejudicial to the reputation of the author of the original work, and tarnishes the reputation of the Trade Marks and each of them.

57.11 Yet further still, the Trade Marks and each of them (in particular when they are included as watermarks as set out above at paragraph 29.4) guarantee to members of the public that the works to which they are affixed are genuine photographs or pieces of footage. This guarantee will be eroded by the Defendant's use complained of herein, in particular, by affixing the sign GETTY IMAGES or ISTOCK to the synthetic output as a watermark members of the public will no longer be able to rely upon the Trade Marks as guaranteeing the authenticity of the works to which they are affixed, thus tarnishing their reputation".

487. My understanding of this pleading is that unfair advantage is pleaded at paragraph 57.7, dilution of distinctiveness is pleaded at paragraph 57.8 (in which the plea of change of economic behaviour is expressly made) and tarnishment is pleaded at paragraphs 57.9-57.11. It is unclear how the allegation of tarnishment involving the appearance of watermarks* on pornography, violent imagery or propaganda could possibly apply to the ISTOCK Mark – there is no evidence of an iStock watermark* appearing on any images of this type and, accordingly, I dismiss any such claim.

Change in Economic Behaviour

488. I agree with Stability that, as is now common ground, the only express plea of change of economic behaviour is set out in paragraph 57.8 of the Particulars of Claim and is concerned with the allegation of detriment to distinctive character/dilution. This is not in fact a plea that there has already been a change in economic behaviour, but rather a plea that there is a serious likelihood of such a change.

489. In closing, Getty Images focused most of its fire power on the allegation of detriment to reputation, arguing that a change in economic behaviour can clearly be deduced from the appearance of the Signs on pornography, violent imagery and propaganda, as well as distorted or manipulated images and synthetic (or fake) images. It submitted that affixing a claimant's mark to a pornographic image is the archetypal case of tarnishment and that, similarly, attributing distorted or manipulated images to Getty Images is

damaging to the reputation it has for licensing high quality, genuine images reporting on editorial content.

490. Getty Images did not refer (either in opening or closing) to its pleaded case on change of economic behaviour with, I consider, good reason. The case as pleaded is unsustainable. It presupposes that the average consumer will wish to use the Model in order to try to circumvent the need to pay for Getty Content. However, the pleading does not say how this will be done. One might postulate in light of the evidence in this case that it could be done by using Getty Images' captions (as Ms Varty did), but I have already held that there is no evidence of this occurring by users of Stable Diffusion in the real world.
491. In any event, even assuming that a real life consumer were to take this step, this is a trade mark claim and so the resulting image would have to bear a watermark*. In other words, as Stability correctly points out, on this hypothesis, after all his effort, the crafty consumer seeking to avoid a license fee, ends up with what he was seeking to avoid: another watermarked* image – the generated image (given the images in Annex 8H) being almost certainly an imperfect and distorted facsimile of the original image. Aside from the fact that it is impossible to see why the average consumer would change his economic behaviour in order to generate a second unusable image with a watermark*, this case is predicated upon the consumer trying to reproduce an image in circumstances where he cannot possibly think anything other than that the appearance of the watermark* is a consequence of his deliberate attempts to reproduce the Getty Images Content. Such a consumer will certainly not regard the appearance of the watermark* as an indicator of origin by Stability.
492. As Stability points out there is no evidence whatever in the case to support the pattern of behaviour pleaded by Getty Images and no basis on which I could properly deduce this change in economic behaviour. The Getty Images' pleading at 57.8 is, in my judgment, no more than pure supposition. On proper analysis it has no foundation in reality and I reject it.
493. Strictly, as it seems to me, Stability is right to say that beyond what is pleaded in paragraph 57.8 of the Particulars of Claim, there is no other pleaded case on change in economic behaviour. Ms Lane sought to address this by pointing out that the court is entitled to make a deduction as to the likelihood of a change in economic behaviour. This is correct for the purposes of determining the issue and, although the pleading could and should have addressed the point, I am most reluctant to determine the case under section 10(3) TMA purely on a pleading point. Stability did not expressly deny the allegations of unfair advantage and tarnishment in its Defence on the grounds that there was no express plea of change in economic behaviour, choosing instead to focus on its arguments as to use in the course of trade. Furthermore, Stability has had the opportunity to cross examine Getty Images' witnesses on this issue, including (in particular) Mr Stanley, just as it has had the opportunity to make detailed submissions on each of the heads of injury alleged. The assessment for the court is a global assessment. In the circumstances, I can see no real prejudice to Stability in the court considering whether it would be appropriate to infer or deduce a change in economic behaviour (in so far as such change has now been articulated by Getty Images) in respect of any of the allegations of injury under section 10(3) TMA.

494. I turn to look in more detail at Getty Images' case on each of the three alleged heads of injury.

Detriment to Distinctive Character

495. I have already dealt with Getty Images pleaded case on dilution, which does allege a change in economic behaviour, albeit not of a type that is in any sense realistic. As I understood Getty Images' case on dilution in closing, however, it was focused on the argument that (i) the Marks are inherently distinctive such that they will be brought to mind inherently and strongly when the average consumer encounters the Sign; and (ii) the use of the Signs will "inevitably weaken the Mark's ability to identify the goods and services for which they are registered as coming from Getty Images, as **there has been, and will continue to be, a proliferation of synthetic image output images bearing the marks**" (emphasis added).

496. However, this case appears to me to be both unusual and problematic. It is now three years since the Model was first made available and although Getty Images has been able to produce evidence to satisfy me that at least one user in the UK will have encountered a watermark*, it has been able to do so only in respect of v1.x and v2.x of the Model, and even then only by reference to an extremely small number of examples. I have found similarity between Mark and Sign in relation to three examples only. While I accept that various of the other watermarks* on real life images in Annex 8I (and on images in Annex 8H where the prompts have included the words "news photo" or "vector art") would likely give rise to a finding of similarity and potentially also to a "link", I have no means of assessing the scale of the production of watermarks* falling within that category in real life. The case is thus historic and my findings at this trial are, of necessity, extremely limited. Getty Images have not sought to run any case based on probability as to the number of watermarks* that may have been produced in the real world by users in the UK of v1.x and v2.x or (more importantly) the number of actually infringing watermarks* that may have been produced (bearing in mind all that I have said already about the fact sensitive nature of the enquiry).

497. Nevertheless, given that I have found that at least one infringing watermark* has been produced by versions 1.x and 2.x, and given that Stability has no case of *de minimis*, should the court infer that there has been detriment?

498. On the available evidence, there is certainly no basis whatever on which to find (or deduce) that there will "continue to be a proliferation of synthetic output images bearing the Marks". I have dismissed Getty Images' case in relation to the later versions of the Models. The unchallenged evidence of Mr Bandari was that since he joined Stability in August 2023, new models have been produced in the form of SD3.0, SD3.5, Stable Video, Stable 3D, language and audio models. There are no allegations in relation to any of these Models.

499. The highest Getty Images were able to put their case on this in closing was to say that the documentary evidence shows that "SDXL was producing watermarks during the development process". This is of course correct but (i) there is no evidence before the Court of any iStock watermarked* image (on any version of the Model) being produced in real life since January 2023 (the Dreaming Image); and (ii) no evidence before the Court of any Getty Images watermarked* image being produced in real life since shortly prior to the public release of SD XL (March 2023). There are no Reddit

Exchanges dating from later than 2023 and the SalesForce Materials on which Getty Images rely date from 2022. The fact that there is evidence that v2.x remains available for download takes matters no further.

500. In all the circumstances, there is certainly no evidence on which I can deduce that there will be a continuing use of the Signs which will dilute the distinctiveness of the Marks.
501. As to historic use, I must bear in mind that the Getty Images Marks have a strong reputation and distinctive character and that (at least where the watermark* is clearly rendered on a synthetic image) its appearance on a synthetic image will immediately and strongly bring to mind the earlier Mark. In such a case there is a greater likelihood that the use of the Sign is detrimental to the distinctive character of the Mark.
502. Ordinarily, in a more conventional case, these points might go a long way towards a finding of dilution. However, I do not consider such a finding to be appropriate in the novel circumstances of this case.
503. Although Stability pleads no case of *de minimis* and accepts “non-trivial” generation of watermarks* by v1.x, the burden of establishing dilution rests with Getty Images and I cannot see that they have satisfied that burden in the very particular circumstances of this case. Their pleaded case on change in economic behaviour is unsustainable and their alternative case, articulated only at trial, requires a finding or deduction that there has been a “proliferation of synthetic images” which satisfy all the conditions referred to above for infringement under section 10(3) and which have led to an adverse effect on the origin function of the Marks.
504. I cannot make such a finding or deduction on the available evidence given the fact-sensitive nature of the enquiry in relation to every generated watermark* and the lack of any probabilistic analysis as to the number of watermarks* that are likely to have been generated in real life by users in the UK. The mere presence of a similar sign on the market is insufficient to establish detriment to the uniqueness and so the distinctive character of the mark. I have no means of determining the scale of the generation of infringing watermarks* by v1.x or v2.1 and not a scrap of evidence of any change in economic behaviour by the users of those Models. Accordingly, having regard to all the circumstances of the case, I consider that there is no evidence on which I could properly deduce that the distinctive character of the Marks has been diminished. Getty Images’ submissions amount to no more than an invitation to engage in supposition, which I decline to do.
505. Getty Images’ case of detriment to distinctive character/dilution under section 10(3) TMA fails for all the reasons set out above.

Detriment to Reputation/Tarnishment

506. The Experts were in agreement that:
 - i) NSFW and violent content is contained in the LAION-5B training dataset. Professor Farid points out in his report that LAION acknowledged that “current automated filtering techniques are far from perfect: harmful images are likely to pass”; and

ii) models trained on such content are in turn capable of producing this content as it is effectively encoded in the model weights. The Model Card for v1.4 expressly confirms that:

“The model was trained on a large-scale dataset LAION-5B which contains adult material and is not fit for product use without additional safety mechanisms and considerations”.

507. As Professor Farid explained, it is possible to filter for such content, both at the training stage (by excluding NSFW images and captions from the training datasets) and at the inference stage. As I have already discussed earlier in this judgment, similar filtering can also be done for watermarked* images.

508. The Experts agree that they are not aware of any public materials which address any specific filtering to remove violent, NSFW images from the datasets used to train and develop v1.x. However, Professor Farid explains in his report that by the time of the release of v1.4, a Safety Checker had been put in place to check model outputs and a prompt filtering step had been applied at the inference stage. Professor Farid explains, and I find, that prompt filtering can be performed reasonably reliably through an API “but is nearly impossible to enforce in open source software because a user can simply remove these safeguards”. I did not understand Professor Brox to disagree.

509. It is clear that (at least by the time of the release of v2.0), Stability had taken steps to filter NSFW content out of the training dataset. The Model Card for v2.0 says:

“The model was trained on a subset of the large-scale dataset LAION-5B, which contains adult, violent and sexual content. To partially mitigate this, we have filtered the dataset using LAION’s NSFW detector”.

The Hugging Face page for v2.0 makes reference to this filter, recording that:

“The training data is further filtered using LAION’s NSFW detector, with a “p_unsafe” score of 0.1 (conservative)”.

LAION’s website reports the accuracy of the filter (or classifier, as Professor Farid describes it) to be 96%, subsequently labelling 2.9% of the dataset as ‘unsafe’. Professor Farid explains that he does not know whether this figure accurately corresponds to the true number of unsafe images in this dataset.

510. Professor Brox confirmed in cross examination, that he would understand this evidence to mean that Stability had not managed fully to mitigate the effect of the inclusion of NSFW in the training dataset, such that there would remain potential for the Model to generate such content. The Experts agree that they are not aware of any public materials which address any express filtering of v2.x of the Model to remove violent images.

511. It was clear from the evidence of both Mr Auerhahn and Mr Bandari that they were aware that there was a problem with the generation of NSFW images by the Model. The following exchange took place during the cross examination of Mr Auerhahn:

“Q. Mr. Auerhahn, it is right, is it not, that the model can produce pornographic images?

A. Yes. It depends on which model we are talking about and what your definition of "pornography" is, but I think one could say that.

Q. Yes. The models we are talking about, just to be clear, are Stable Diffusion v1, v2 and XL.

A. Right. By the time it got to v2 and XL it was definitely not as good at doing that and would not do it as often, and many times took most of that stuff out, but it is still possible.

Q. It did do it occasionally?

A. Yes”.

512. Mr Bandari confirmed in cross examination that he had been involved in a Slack Chat exchange in April 2024 following contact from a professor who had informed Stability that “it was very easy to get around [its] safety filter and get really violative results”. The means by which this can be achieved is later explained in the Slack Chat by Javad Azimi who says this:

[Redacted]

Javad Azimi observes that the “image classifier...needs some improvement”.

513. In the same Chat, Mr Konstantinov refers to this as a problem with DreamStudio XL generating NSFW images: “Generating NSFW with DreamStudio (SDXL) doesn’t seem to be a big deal, even if you don’t use fancy pipelines”. He identifies a different means of achieving NSFW images which appears to involve using the image to image function on the model: [Redacted]

514. In the same Chat, Ella Irwin says that she is worried that “Daria and Misha were easily able to producing (*sic*) NSFW content using [Dream]Studio in like two seconds. Why is that? It’s my understanding that we have had problems with NSFW content and pretty serious NSFW content with DreamStudio in the past. Did the filters stop working?”. Mr Konstantinov says that Stability would need “to develop and apply filters both for DreamStudio and for our future assistant”. Mr Bandari agreed in cross examination that this situation might well be worse where users had downloaded the Model because Stability could not filter the prompts used locally. I note that in the same Chat, Todd Elvers points out that “Safety filters for API calls from DreamStudio should be same as SD XL 1.0 and SD1.6” and he provides a link. Later in the Chat Mr Bandari says that “the issue is not isolated to DreamStudio”, by which he appears to mean it applies to all APIs. He goes on to observe that “we will have to find ways that we can reproduce easily and see if we can mitigate it”. The Chat continues to discuss how the problem can be addressed. Mr Bandari was not questioned about the detail of the means by which the filters could be by-passed or the possibility that this would be done by real world users. Mr Konstantinov was also not asked about this.

515. Ms Hodesdon accepted in cross examination that the Model can produce outputs that are NSFW but explained that Stability was able to exercise control over outputs “to ensure that we are not serving that content back to end users”. She confirmed that “pretty robust” filters were used for this purpose on Stability’s server, that these filters had improved over time, that “at the moment they are subject to lots of testing”, and that Stability can “control the sensitivity of the filters”. She confirmed that this evidence did not apply to source code and model weights downloaded from GitHub and Hugging Face – evidence consistent with Professor Farid’s views.

516. Getty Images say that given that the generation of NSFW images was plainly a problem, and given that there is evidence that the generation of watermarked* images was also a problem, “the combination of the two is also likely”. Thus it contends that any association with such content will tarnish the reputation of the Marks. Getty Images rely on the fact that, during the Getty Watermark Experiments, images of Miley Cyrus were generated which they say are pornographic and which bear Getty Images watermarks*. These are referred to in paragraph 57.9 of the Particulars of Claim as appearing at pages 20 and 100 of Annex 8H and in a confidential exhibit to Mr Stanley’s statement:

- i) Page 20 of Annex 8H shows an image generated by Ms Varty using the prompt: “Miley Cyrus at American Awards, news photo”. A woman appearing to be Miley Cyrus appears three times in the same image on the red carpet wearing different clothing and, in one case, exposing her breast. This image was generated by v1.2.
- ii) Page 100 of Annex 8H shows an image generated by Ms Varty using the prompt: “Miley Cyrus, performs onstage during the 2008 American Music Award, news photo”. A woman appearing to be Miley Cyrus appears twice in the image, again exposing her breasts. This image was generated by v2.1.
- iii) Mr Stanley’s exhibit includes the two images referred to above together with 17 other images of Miley Cyrus exposing various parts of her body and (in one case) appearing to have an arm extending out of her crotch. These were all generated by Ms Varty using the two prompts identified above. They were generated by v1.2, 1.4 and 2.1 of the Model.

517. It is Getty Images’ case that the content of these images, specifically the depiction of nudity and the particular “poses and actions of the persons depicted” renders them NSFW. Mr Stanley confirmed in his evidence that he believed that these images “could be considered pornographic”. The Court of Appeal ([2025] EWCA Civ 749) characterised them as “artificial images of celebrities in states of undress” and took the view (at [28]) that “[t]hey therefore fall within the broad remit of pornography”.

518. Stability says that, on the balance of probabilities, it is “vanishingly unlikely” that any pornographic images with watermarks* will have occurred in the UK. It points out that none of the images relied upon by Getty Images was generated through DreamStudio or via the Developer Platform (in respect of which filtering is possible and is done) and that there is no evidence whatever of such NSFW images with watermarks* being produced in real life, notwithstanding Getty Images’ investigations of social media and its millions of interactions with Getty Images customers since this litigation began.

519. Stability points out that at least 1,000 images of Miley Cyrus were generated during the litigation of which only 19 include watermarks*. It also points out that the scenarios identified in the Slack Chat all involve deliberate misuse of the Model and are no more than “internal stress testing for preventative measures” which say “nothing about user interactions”. It relies again on what Getty Images said to the Court of Appeal to the effect that it is users who choose to generate pornography, violent images and propaganda – their actions have nothing whatever to do with Stability.

520. On balance, I agree with Stability that this is not a case in which I can properly find damage to reputation in relation to:

- i) the production of images bearing the iStock watermark*. There is no evidence whatever of any pornographic or other NSFW images being produced bearing this watermark*.
- ii) the production of images bearing the Getty Images watermark* by users accessing v1.x via the Developer Platform/API, or accessing v1.4 via DreamStudio (the only circumstances in which Stability would have responsibility for watermarks* generated by v1.x of the Model). There is no evidence of any pornographic or other NSFW images being produced in these use-cases in real life, much less of any such images bearing watermarks*. Ms Hodesdon said that she was not aware of any case in which a user had received an output from Stability’s inference server that “they have complained to be Not Safe For Work” – it is difficult to believe that if children looking to generate an image of their favourite popstars have accidentally been generating pornographic images (as posited by Getty Images), there would not have been a barrage of complaints.

521. Although I have found that evidence of images bearing watermarks* may be difficult to come by owing to the fact that many consumers who do not want those images will simply discard them, it seems to me that the position is rather different when one is talking about potentially shocking or disturbing pornographic images bearing Getty Images or iStock watermarks*. If this had been happening in real life in response to prompts that were not designed to generate a pornographic image (including, for example the words “news photo”) I would have expected to see evidence in social media or in Getty Images’ SalesForce Materials of that. I would also have expected to see evidence in Stability’s internal Chats addressing the potential for pornography to appear together with watermarks*. However, there is no mention of anyone coming across a pornographic image bearing a watermark* in the cases referred to above (or indeed in any real life instance), notwithstanding that extensive “stress-testing” on the production of pornographic images appears to have been taking place.

522. Nonetheless, it is true that a real world user could have used the prompt “news photo” in relation to v2.1 of the Model in conjunction with other text (perhaps using the name of a female popstar), and (having done so) could have generated a NSFW image with a Getty Images watermark*. The Getty Watermark Experiments confirm that this can happen in relation to v1.x and v2.1 of the Model (at least insofar as the download mechanism is concerned).

523. Against that background, the question for the court is whether, absent any evidence of “real life” generation of pornographic images bearing watermarks*, the Miley Cyrus

images (generated by Getty Images during their experiments and with their consent), together with the Expert evidence to which I have referred, are sufficient to enable me to deduce damage to reputation.

524. In my judgment they are not, for the following reasons:

- i) There is not a scrap of evidence that watermarks* have in fact appeared on pornographic images in the real world at any time and in relation to any Model. There are no examples of watermarks* appearing on NSFW images other than of Miley Cyrus and there is no evidence to explain why this might have happened in relation to the Miley Cyrus image (although Mr Stanley confirmed in his cross examination that he had recently become aware that there were NSFW images of Miley Cyrus available on the Getty Images Websites (albeit as I understand it these are of course real images and come with appropriate warnings)). This may provide the explanation for the appearance of watermarks* on the Miley Cyrus images, but Mr Stanley was unable to assist on this point.
- ii) It appears that captions including Miley Cyrus' name were included as the first three entries in the list indexed to the US complaint from which Ms Varty began the Getty Watermark Experiments, but there is no evidence from the individual who drafted the list to explain why Miley Cyrus was chosen. I cannot begin to say whether the Miley Cyrus images might be representative of other images that could have been generated in the real world using the caption "news prompt" together with the name of a celebrity, or not. In light of Mr Stanley's evidence, the answer may depend upon whether there are any other (and if so how many) NSFW (pornographic) images on the Getty Images Websites and whether the Models (or any of them) were trained on these images. However, I have no evidence to assist on this.
- iii) Although I accept Mr Stanley's evidence (i) that the Miley Cyrus images are an example of the "real risk" posed to the Getty Images' brand by the potential for the Model to generate inappropriate NSFW imagery (as qualified in his evidence to explain that he was describing "severity" and not "probability"); and (ii) that even if one such image has been generated in the real world that could have been "problematic" to the Getty Images brand, this seems to me to be too theoretical in the circumstances of this case. The court simply has no means of determining whether anything similar has in fact been generated in the real world – Mr Stanley was unable to produce any further examples. Professor Farid said no more than that Models trained on NSFW are capable of producing this content in their synthetic outputs. He expressed no view on whether such outputs would bear watermarks*.
- iv) That a problem with NSFW images was picked up by a professor and that Stability recognised that there was an issue internally, does not appear to me to be sufficient to support a deduction that pornographic images bearing watermarks* have been generated in response to innocuous prompts in the real world. As Stability points out, the type of prompts apparently required to "get around" the existing filters used in DreamStudio and on the API are somewhat "niche" and provide no support for the proposition that the average consumer

will obtain such images by putting in random text prompts, much less that if he does obtain such an image it will also bear a watermark*.

- v) While there is no doubt NSFW images may be produced by the Models (as confirmed by the Experts), the evidence suggests that it is much more likely that this will occur where positive efforts are made to produce them – probably by users who have downloaded the Model, thereby enabling them to get around the “robust filters” applied by Stability to outputs generated on its own computer infrastructure for users accessing the Model via DreamStudio and the Developer Platform. In such circumstances, and consistent with what I have said earlier in this judgment, I agree with Stability that there could be no rational basis on which the average consumer would understand any watermark* that happened to be generated in the course of such behaviour as a commercial communication about the origin of the image or the service they were using or indeed as a message that Getty Images condoned pornography, as opposed to “an artefact associated with their own peccadillos” (as Stability put it). These submissions are echoed by the submissions made by Getty Images to the Court of Appeal about the generation of CSAM images.
- vi) In so far as Getty Images’ allegation seeks to rely upon violent images and images of propaganda, its experiments did not even produce any examples. In closing, Ms Lane explained that she was relying on the evidence contained in the Model Cards together with the Donald Trump Image. However, that image was generated by v1.5, in respect of which there is no case and there is simply no other evidence whatsoever of images containing violence or propaganda bearing watermarks* being generated in the real world.
- vii) It is not enough, as it seems to me, for Getty Images to assert that because the Models are obviously capable of generating NSFW images, those images will on occasions have borne watermarks*. Given the lack of any probabilistic case as to the incidence of watermarks* generally, or the potential for them to appear on pornographic images, together with the total lack of any real life evidence (other than the Miley Cyrus images) I cannot see that it would be appropriate to deduce that there must have been damage to reputation. That would be pure supposition.
- viii) Of course, as Getty Images pointed out, the use of a Sign in an unpleasant, obscene or degrading context, incompatible with the reputation and image of the Mark, will frequently give rise to tarnishment. It is unsurprising that Mr Stanley expressed concern. If watermarks* have appeared on pornographic (or other NSFW) outputs inadvertently generated by UK users of the Model then that would plainly provide good and persuasive evidence of tarnishment. On the available evidence before the court and in all the circumstances of this case, however, there is no such evidence and I cannot find tarnishment.
- ix) There is no basis on which I could find, or deduce, a change in economic behaviour.

525. In its pleading at paragraph 57.10, Getty Images also seek to rely upon the damage to reputation that would be caused by distorted or manipulated images. Curiously, however, 57.10 appears to be parasitic on Getty Images’ case on primary infringement

of copyright, which they have abandoned. The express premise is that the watermark* is accompanying a copied underlying image and that it is prejudicial to the reputation of the author of that image as well as tarnishing the reputation of the Marks.

526. In closing, Getty Images sought to address this allegation together with the subsequent allegation at paragraph 57.11 in relation to “non-genuine photographs”. It did not explain how I was to address the allegation that watermarks* appeared on “copied” images or how this plea could sensibly survive the abandonment of the Outputs Claim. Aside from the SOCI images (produced with the consent of Getty Images) there is no evidence that synthetic images have been generated by the Models which distort an underlying image from which it has been copied.
527. At 57.11, Getty Images focuses on “non-genuine photographs”. Ms Cameron explained Getty Images’ concern about “deepfakes” – i.e. images where the content has been digitally altered and used to spread false and misleading information: “We want customers to know that any editorial content that they source from Getty Images is what we say it is. It is not manipulated. It has not been altered. It is authentic.” Ms Gagliano gave similar evidence.
528. As I understood this evidence it was directed at the editorial side of Getty Images’ business – i.e. photographs. It cannot relate to the ISTOCK Marks and Getty Images’ pleading is expressly limited to “genuine photographs or pieces of footage”.
529. There is no evidence whatever of deepfakes being created by the Models and used in the real world to spread misinformation – this is no more than supposition.
530. Finally, Getty Images emphasise Mr Stanley’s evidence as to the damage to Getty Images’ reputation that may be caused by the poor quality of the images synthesised by Stable Diffusion. This arises in the context of Getty Images’ claim that Getty Images watermarks have a reputation as guarantees of “genuine photographs or pieces of footage”.
531. Annex 8H shows that synthetic images may be generated (bearing watermarks*) in conjunction with the use of the words “news photo” and comparatively simple prompts which appear at best of low quality and at worst significantly distorted. I accept Mr Stanley’s evidence that the generation of low quality “non-genuine” photographs is a concern for Getty Images “because we pride ourselves on the quality of our photography and our product. If people were to think that these images came from, or were authorised by or associated with Getty Images, the fact that they are of such low quality would damage our reputation for producing high quality content”. Ms Cameron’s evidence was to similar effect. Two examples appear below, generated using v1.2 with the text captions “news photo of MTV VMAs” and “Obama, news photo”:



532. The watermarks* in these images are plainly “off” (appearing twice on each image and being extremely distorted) as are the images on which they have been affixed. However I am prepared to accept that depending on the appearance of the watermarks* and the image, there may be scope for the Model (versions 1.x and 2.x) to produce distorted images which appear photo-realistic and which are not so “off” that the average consumer would appreciate there could be no link with Getty Images. If that had happened in relation to outputs generated by relevant versions of the Model in real life then there might have been scope for a tarnishment claim.

533. The trouble is, however, that because each watermark* and each image will be different, and because I cannot say how often this will have happened in real life, I do not consider that I am in a position to find tarnishment on this basis in the circumstances of this case. Until after the circulation of the draft judgment, I was not asked to make any specific findings about the images that Getty Images produced using the “news photo” prompt or the watermarks* appearing on them, and accordingly I have declined to do so for reasons set out above. Furthermore, there is no evidence on which I could make any findings as to the frequency with which users will input prompts which include the words “news photo” or the probabilities of such prompts generating images that bear watermarks* which are sufficiently clear that a link will be made with Getty Images. This is (again) a function of the particular difficulties that arise in a case where no two manifestations of the Sign, or the image on which it appears, are the same. Once again,

there is also no evidence whatever of a change in economic behaviour and no basis for deducing any such change.

534. Getty Images' case of detriment to reputation/tarnishment fails for all the reasons set out above.

Unfair Advantage

535. Getty Images' pleading of unfair advantage is wholly unparticularised. It presupposes that Stability was intending to take advantage of the Claimants' reputation in using the Sign, alternatively that that was the objective effect of its use. I bear in mind the importance of undertaking a global assessment on this question but on the facts of this case, this claim is unsustainable. It is true that the Models were trained on datasets which included Getty Images Content with watermarks. But it is clear from the evidence that the presence of watermarks* on synthetic images is undesirable and makes them unsuitable for use. Users of Models which produce images with watermarks* are, as I have held, likely to discard such images.

536. There is no basis on which I can infer that Stability intended watermarks* to appear or that it wished to take advantage of Getty Images' reputation by reason of that appearance, and none was suggested. Indeed the evidence shows that from v2.x, Stability sought to take steps to try to filter out watermarks*: Getty Images' pleaded case is that v2.x was trained on LAION-A which was created by the DeepFloyd team based on the filtering of LAION 5-B. Mr Konstantinov confirms that this involved a watermark filter and I agree with Stability that the fact that Getty Images was unable to generate any iStock watermarks* using v2.x during its experiments suggests that this was at least partially successful.

537. I can see no basis whatever on which it can sensibly be argued that Stability intended to gain any unfair advantage from watermarks* appearing on synthetic images generated by the Models and nor do I consider that to be the objective effect of the appearance of Signs on synthetic images. Notwithstanding the extensive reputation of the Marks and their distinctive character, it is very difficult to see how that reputation and character will be transferred to Stability by reason of their use of the Signs in circumstances where the average consumer will not wish to use the images generated.

Conclusion on section 10(3) Infringement

538. For all the reasons set out above, Getty Images' claim under section 10(3) TMA fails.

(G) PASSING OFF

539. It is common ground that the applicable legal principles were recently set out by the Court of Appeal in *Lidl* at [27]-[35]. As Arnold LJ there observed (at [29]), the most comprehensive statement of the law remains that of Lord Diplock in *Erven Warnink BV v J Townend & Sons (Hull) Ltd* [1979] AC 731 at 742:

"My Lords, *A. G. Spalding & Bros. v. A. W. Gamage Ltd.*, 84 L.J.Ch. 449 and the later cases make it possible to identify five characteristics which must be present in order to create a valid cause of action for passing off: (1) a misrepresentation (2) made

by a trader in the course of trade, (3) to prospective customers of his or ultimate consumers of goods or services supplied by him, (4) which is calculated to injure the business or goodwill of another trader (in the sense that this is a reasonably foreseeable consequence) and (5) which causes actual damage to a business or goodwill of the trader by whom the action is brought or (in a *quia timet* action) will probably do so.”

540. By closing submissions, it was Getty Images’ case that the claim for passing off would ultimately “stand or fall” with the Trade Mark Infringement Claim and they relied upon the evidence to which I have already referred in that context.
541. However, as Stability points out, there is at least an argument that passing off does not extend to post-sale confusion – the only exception in the authorities being where the party purchasing the goods does so in order to make a further misrepresentation and the trader intentionally trades off the back of this intention. See the analysis of *Freddy SpA v Hugz* [2020] EWHC 3032 (IPEC) in *Thom Browne v Adidas* [2024] EWHC 2990 (Ch) at [755]-[758]. The court did not need to decide the point in that case, but on any view, the question there articulated may potentially be relevant here. The tort with which I am concerned only crystallises after the user has made the decision to download the Stable Diffusion model weights and source code, or to sign up for DreamStudio or to use the Developer Platform/API.
542. Getty Images did not address me on this point but simply asserted that if the case on sections 10(2) and/or 10(3) TMA succeeded, then they had a legitimate claim of misrepresentation which had deceived or was likely to deceive the public and which had caused damage. I am not so sure that this analysis is correct on the facts of this case, but where the point has not been properly argued before me I decline to address it. Getty Images has succeeded in part in its claims under sections 10(1) and 10(2) TMA and I do not presently see that its claim of passing off adds anything to my findings on those claims. I gave the parties the opportunity to make further submissions on this point upon circulating this judgment in draft, but neither party sought to do so and indeed it was Stability’s position that the court ought not to entertain further substantive argument on the issue after conclusion of the trial. In circumstances where neither party invites me to consider the point further, there is no need for me to do so.

(H) THE SECONDARY INFRINGEMENT CLAIM

543. The CDPA differentiates between primary and secondary acts of infringement. Secondary acts of infringement are broadly addressed to downstream dealings or involvement, as opposed to acts which originate reproductions of copyright works. Getty Images contend that Stability has committed two related acts of secondary infringement of copyright contrary to sections 22 and 23 CDPA.
544. Getty Images’ case has narrowed by reason of the abandonment of the Training and Development Claim. Absent that claim, Getty Images must now focus solely on a claim that Stable Diffusion is an “infringing copy” because it has been imported into the UK and its making in the UK would have constituted an infringement of the copyright in the Copyright Works pursuant to CDPA section 27(3). Getty Images contend that infringement has occurred by reason of the importation of Stable Diffusion (through its downloading in the UK) and its distribution in the course of business via Hugging Face.

The Statutory Framework

545. Section 22 CDPA provides as follows:

“22. Secondary infringement: importing infringing copy.

The copyright in a work is infringed by a person who, without the licence of the copyright owner, imports into the United Kingdom, otherwise than for his private and domestic use, **an article which is, and which he knows or has reason to believe is, an infringing copy of the work**” (emphasis added).

546. Section 23 CDPA is in similar terms:

“23. Secondary infringement: possessing or dealing with infringing copy.

The copyright in a work is infringed by a person who, without the licence of the copyright owner—

- (a) possesses in the course of a business,
- (b) sells or lets for hire, or offers or exposes for sale or hire,
- (c) in the course of a business exhibits in public or distributes, or
- (d) distributes otherwise than in the course of a business to such an extent as to affect prejudicially the owner of the copyright,

an article which is, and which he knows or has reason to believe is, an infringing copy of the work” (emphasis added).

547. “Infringing copy” is defined in section 27 CDPA as follows:

“27. Meaning of “infringing copy”.

- (1) In this Part “infringing copy” in relation to a copyright work, shall be construed in accordance with this section.
- (2) An article is an infringing copy if its making constituted an infringement of the copyright in the work in question.
- (3) An article is also an infringing copy if—
 - (a) it has been or is proposed to be imported into the United Kingdom, and
 - (b) its making in the United Kingdom would have constituted an infringement of the copyright in the work in question, or a breach of an exclusive licence agreement relating to that work.”

The key Issues between the parties

548. The battle lines can be readily and simply identified. Getty Images' case raises two novel issues of law together with a number of related issues of fact. The issues of law are whether Stable Diffusion is capable of being (i) an "article" for the purposes of sections 22 and 23 CDPA and (ii) an "infringing copy" for the purposes of section 27 CDPA.

549. Getty Images submit that these questions are both to be answered in the affirmative. They do not say that the Model itself comprises a reproduction of any Copyright Works, but they submit that the definition of "infringing copy" in section 27(3) CDPA is sufficiently broad to encompass an article (including an intangible article) whose creation or "making" involved copyright infringement. They point out that it is common ground that the training of the Model involved the reproduction (by means of storage) of the Copyright Works both locally and in cloud computing resources. They contend that the "article" with which the court is concerned is the model weights and that the "making" (or optimization) of the model weights requires (as is common ground) their repeated exposure to the training data, exposure which "fundamentally alters" their composition. They say that if this "making" had been done in the United Kingdom it would have constituted an infringement for the purposes of section 27(3) CDPA.

550. Stability contends, however:

- i) that the term "article" in the CDPA is properly limited to tangible objects. It says that intangible or abstract information, such as the Stable Diffusion Models in issue which are distributed over the internet in intangible form, are not articles. Stability contends that this construction is clear from the wider context of the CDPA and (if necessary) a reading of Hansard.
- ii) that the term "infringing copy" cannot apply to Stable Diffusion in circumstances where it was trained on Copyright Works in the United States, copies of those works were never present within the Model and where the Model has "never had the slightest acquaintance with a UK Copyright Work". Stability submits that the interpretation for which Getty Images contend is "exorbitant" because the act of training the model weights ultimately did not involve storing or reproducing the images in those weights. Again, Stability points to the wider context of the CDPA in support of its interpretation.

551. The issues of fact to be determined by the court (as identified by the parties) boil down to (i) whether Stable Diffusion is an "infringing copy" as a matter of fact; (ii) whether Stability had knowledge or reason to believe that Stable Diffusion is an infringing copy; and (iii) whether Stability has committed secondary acts of copyright infringement contrary to sections 22 and 23 CDPA.

Preliminary points about the nature of Stable Diffusion

552. I have already set out at some length in the judgment my understanding as to the nature of an AI model such as Stable Diffusion, together with information as to how it was trained and why it produces images bearing watermarks*. At the outset of the judgment

I recorded that (as is agreed by the Experts) Stable Diffusion does not itself store the data on which it was trained.

553. For present purposes, however, it is important that I set out in rather more detail the evidence as to how a model such as Stable Diffusion is trained and how it produces images. In my judgment it is not possible to determine whether Stable Diffusion is capable of being an infringing copy without a clear understanding of what Stable Diffusion actually is.
554. It was Professor Brox's unchallenged evidence in his report (which I accept) that:

“8.36...in order for a diffusion model to successfully generate new images, that model must learn patterns in the existing training data so that it can generate entirely new content without reference to that training data.

8.37 Rather than storing their training data, diffusion models learn the statistics of patterns which are associated with certain concepts found in the text labels applied to their training data, i.e. they learn a probability distribution associated with certain concepts. This process of learning the statistics of the data is a desired characteristic of the model and allows the model to generate new images by sampling from the distribution.

...

8.40 ...For models such as Stable Diffusion, trained on very large datasets, it is simply not possible for the models to encode and store their training data as a formula.... It is impossible to store all training images in the weights. This can be seen by way of a simple (example) calculation. As I explained in paragraph 6.28 above, the LAION-5B dataset is around 220TB when downloaded. In contrast, the model weights for Stable Diffusion 1.1-1.4 can be downloaded as a 3.44GB binary file. The model weights are therefore around five orders of magnitude smaller than a dataset which was used in training those weights”.

555. This evidence is consistent with the agreement in the Expert Joint Statement that “the model weights do not directly store the pixel values associated with billions of training images” – i.e. digital images each consisting of what Professor Farid describes in his report as “an array of pixels”. During training, images are converted from pixel space into “**latent space**” using an **autoencoder**. Latent space is a compressed, representative form of the pixel space image that is more memory and computationally efficient. As explained in the Agreed Technical Primer:

“During training, each image is encoded into a latent representation to which a random amount of noise is added. The network weights are optimized to predict this noise [following a standard approach known as ‘stochastic **gradient descent**’] so that the network can recreate the content that is destroyed by the additive noise. Besides the noisy images, this optimization also

takes into account the paired text description associated with each image.

The optimization of the weights follows a standard approach known as ‘stochastic gradient descent’. Given the status of the weights, the locally best change of these weights (the gradient) is computed to optimize the noise prediction for the training samples at hand. The gradients of training samples are accumulated over a small subset of the training dataset: the **batch**. Its size is referred to as the batch size. After taking a step in the direction of this accumulated gradient, the training procedure selects a new batch for the next **step**, and so on until all the images in the training data are processed, defined to be a single ‘**epoch**’. This entire process may be repeated for more than one epoch

A single training image only has a limited impact on how the parameters of the network will be changed. It is competing with the gradients of the other samples in the batch, which may not necessarily agree on the direction of the gradient. Because each batch is a different subset of the full training dataset, the gradients in each step will not include a training image seen in a previous batch (unless that image is itself replicated in the training images). Rather, it will only be seen again after all training images have been seen once (an epoch), and going over the training images is repeated

The network training makes the most efficient progress when many training samples agree on a local or global pattern. In this case, some of the weights are changed consistently by all agreeing samples”.

556. As Professor Farid explains in his evidence, typically at this development stage, the annotated training dataset will be stored locally on a desktop or laptop where the initial development is performed. Indeed Stability confirms in its Defence that prior to training copies of each image in the dataset were downloaded and stored on Amazon S3 on the AWS Cluster and that, during each iteration of training, images were retrieved by cloud computing devices located within the AWS Cluster and temporary copies of the images were made in the Video Random Access Memory (“**VRAM**”) of the GPUs performing the training on the AWS Cluster. It is now accepted that this did not take place in the UK.
557. Once the network is trained, the process of inference is conducted without the need for any training data to be used.
558. The Experts agree that, while Stable Diffusion can and does produce images that are “distinct from the training examples”, it can also produce “images that are nearly identical (a memorized image)”, and that it can produce “images that are derived from a training image either in part or in whole (a derivative)”. The Agreed Technical Primer records that:

“The network’s weights are optimised on the training data, but its goal is to perform well on previously unseen data. In the context of Stable Diffusion, unseen data means new random noise patterns and/or new text inputs. To work reasonably on such new data, the network must be able to ‘generalise’: to recognise and understand the general patterns and rules in the training data and be able to apply them in a different context.

If a network has been trained for too long on the same training data or an insufficiently diverse training data, it can be prone to ‘overfitting’. Overfitting occurs when the network uses its weights or part of its weights to memorize the individual training images rather than representing a large set of training images jointly with these weights. Overfitting is characterised by small errors on the training data, but a high error rate on new, unseen data. Overfitting is an undesired feature in machine learning, which engineers try to avoid.

Deep networks can both generalize and memorize at the same time. In such case, the network uses most of its weights to represent general patterns in the data, but uses some part of its weights to memorize individual patterns. The presumed primary cause for memorization is duplication of training data, either by explicit duplication or by training the network for too many epochs, in conjunction with patterns that cannot be easily represented together with other patterns in the dataset – so-called “outliers”.

559. However, notwithstanding this evidence about memorization, it is important to be absolutely clear that Getty Images do not assert that the various versions of Stable Diffusion (or more accurately, the relevant model weights) include or comprise a reproduction of any Copyright Work and nor do they suggest that any particular Copyright Work has been prioritised in the training of the Model. There is no evidence of any Copyright Work having been “memorized” by the Model by reason of the Model having been over-exposed to that work and no evidence of any image having been derived from a Copyright Work.
560. Getty Images’ case is that they do not need to show any of this; they say it is enough that the making of the model weights (had it been carried out within the UK) would have constituted an infringement of copyright.
561. Before turning to deal with construction of the statute, I observe that there was considerable argument during submissions about various factual “analogies”, usually identified by Stability. Although designed to be helpful, these were invariably controversial and, for that reason, I do not find them of any real assistance and shall not refer to them again. It seems to me that I must keep a firm eye on the facts set out above and not get distracted by the consideration of different factual scenarios.

The Approach to Statutory Construction

562. The general approach to statutory construction was recently summarised by the Privy Council in *Al-Thani v Al-Thani* [2025] UKPC 35, [2025] 1 BCLC 473 at [23]-[28] and I did not understand it to be in issue:

"23. Turning to the common law, in *R (Quintavalle) v Secretary of State for Health* [2003] UKHL 13; [2003] 2 AC 687, Lord Bingham of Cornhill stated (para 8):

"The basic task of the court is to ascertain and give effect to the true meaning of what Parliament has said in the enactment to be construed."

This is not a licence simply to interpret literally the particular provision and neglect the purpose which the legislature intended to achieve when it enacted the statute.

24. Reading a provision in its statutory and historical context assists in determining the purpose of the provision. Lord Bingham continued in *Quintavalle* (also in para 8):

"Every statute other than a pure consolidating statute is, after all, enacted to make some change, or address some problem, or remove some blemish, or effect some improvement in the national life. The court's task, within the permissible bounds of interpretation, is to give effect to Parliament's purpose. So the controversial provisions should be read in the context of the statute as a whole, and the statute as a whole should be read in the historical context of the situation which led to its enactment."

25. The House of Lords in *R v Secretary of State for the Environment, Transport and the Regions, Ex p Spath Holme Ltd* [2001] 2 AC 349, ("Spath Holme") 396–398 per Lord Nicholls of Birkenhead, and the United Kingdom Supreme Court in *R (O) v Secretary of State for the Home Department* [2022] UKSC 3; [2023] AC 255, ("R(O)") paras 29–31 per Lord Hodge, have also placed emphasis on the importance of interpreting statutory words in their context. In the latter case the court stated (para 29):

"Words and passages in a statute derive their meaning from their context. A phrase or passage must be read in the context of the section as a whole and in the wider context of a relevant group of sections. Other provisions in a statute and the statute as a whole may provide the relevant context. They are the words which Parliament had chosen to enact as an expression of the purpose of

the legislation and are therefore the primary source by which meaning is ascertained."

26. Often reports of advisory committees or explanatory notes which accompany legislation can assist the interpretation of a statutory provision or give guidance on the purpose of the legislation as a whole. But such external aids usually play a secondary role. In *Spath Holme* Lord Nicholls (p 397) explained the constitutional reason for having regard primarily to the statutory words and the statutory context and purpose:

"Citizens, with the assistance of their advisers, are intended to be able to understand parliamentary enactments, so that they can regulate their conduct accordingly. They should be able to rely on what they read in an Act of Parliament."

...

28. In *R(O)* at para 31 the task of the court was described in these terms:

"Statutory interpretation involves an objective assessment of the meaning which a reasonable legislature as a body would be seeking to convey in using the statutory words which are being considered."

...".

563. Getty Images drew my attention in their submissions to the well-established principle that, in general, a statutory provision is "always speaking". In *News Corp UK & Ireland Ltd v Commissioners for His Majesty's Revenue and Customs* [2023] UKSC 7, [2024] AC 89 Lord Hamblin and Lord Burrows JJSC (with whom Lord Hodge DPSC and Lord Kitchen JSC agreed) said this at [27]-[30]:

"27 It is clear that the modern approach to statutory interpretation in English (and UK) law requires the courts to ascertain the meaning of the words used in a statute in the light of their context and the purpose of the statutory provision...

28 Within that modern approach, it is also a well-established principle of statutory interpretation that, in general, a provision is always speaking: see, e.g., *Royal College of Nursing of the United Kingdom v Department of Health and Social Security* ("Royal College of Nursing") [1981] AC 800; *Rv G Ireland* [1998] AC 147, 158-159; *Quintavalle; Owens v Owens* [2018] AC 899 (approving [2017] 4 WLR 74); *Test Claimants in the FII Group Litigation v Revenue and Customs Comrs* [2022] AC 1. See also Greenberg (ed), *Craies on Legislation*, 12th ed (2022), ch 21; and *Bennion, Bailey and Norbury on Statutory Interpretation*, 8th ed (2020), ch 14.

29 What is meant by the always speaking principle is that, as a general rule, a statute should be interpreted taking into account changes that have occurred since the statute was enacted. Those changes may include, for example, technological developments, changes in scientific understanding, changes in social attitudes and changes in the law. Very importantly it does not matter that those changes could not have been reasonably contemplated or foreseen at the time that the provision was enacted. Exceptionally, the always speaking principle will not be applied where it is clear, from the words used in the light of their context and purpose, that the provision is tied to an historic or frozen interpretation. A possible example (referred to by Lord Steyn in *R v Ireland* at [1998] AC 147, 158) is *The Longford* (1889) 14 PD 34 where the word “action” in a statute was held not to be apt to cover an Admiralty action in rem: at the time the statute was passed, the Admiralty Court “was not one of His Majesty’s Courts of Law” (p 37).

30 The great merit of the always speaking principle is that it operates to prevent statutes becoming outdated. It would be unrealistic for Parliament to try to keep most statutes up to date by continually passing amendments to cope with subsequent change”.

564. At paragraph [32], Lords Hamblin and Burrows went on to refer to an extract from the dissenting judgment of Lord Wilberforce in *Royal College of Nursing* [1981] AC 800 at page 822, subsequently approved by the House of Lords in *R (Quintavalle v Secretary of State for Health* [2003] 2 AC 687, in the following terms:

“In interpreting an Act of Parliament it is proper, and indeed necessary, to have regard to the state of affairs existing, and known by Parliament to be existing, at the time. It is a fair presumption that Parliament’s policy or intention is directed to that state of affairs . . . when a new state of affairs, or a fresh set of facts bearing on policy, comes into existence, the courts have to consider whether they fall within the Parliamentary intention. They may be held to do so, if they fall within the same genus of facts as those to which the expressed policy has been formulated. They may also be held to do so if there can be detected a clear purpose in the legislation which can only be fulfilled if the extension is made. How liberally these principles may be applied must depend upon the nature of the enactment, and the strictness or otherwise of the words in which it has been expressed. The courts should be less willing to extend expressed meanings if it is clear that the Act in question was designed to be restrictive or circumscribed in its operation rather than liberal or permissive. They will be much less willing to do so where the new subject matter is different in kind or dimension from that for which the legislation was passed. In any event there is one course which the courts cannot take under the law of this country: they cannot

fill gaps; they cannot by asking the question, ‘What would Parliament have done in this current case—not being one in contemplation—if the facts had been before it?’, attempt themselves to supply the answer, if the answer is not to be found in the terms of the Act itself.”

The meaning of the terms “article” and “infringing copy”

565. The parties dealt with these terms separately in their submissions, but there is a considerable overlap for reasons which will become clear. Each prays in aid the words of the statute, the statutory context, the decision in *Sony Computer Entertainment Inc v Ball* [2004] EWHC 1738 (Ch), [2005] FSR 9 (“**Sony v Ball**”) and the need for a purposive interpretation of the statute in support of their respective interpretations of the words “article” and “infringing copy”. Getty Images emphasise the “always speaking” principle, while Stability invites me to have regard to the record in Hansard of what was said in Parliamentary debates on sections 22 and 23 CDPA in determining the meaning of the term “article”.
566. Getty Images began their submissions with the interpretation of the word “article” and so that is also where I shall begin.

The Interpretation of the word “article”

567. There is no relevant statutory definition of “article”. Getty Images submit that the word “article” is a very general term with a broad meaning and that it does not have to denote a tangible thing. They point to a “news article” or an “article of the constitution”. They also point to the fact that in each of sections 22 and 23 CDPA the term “article” is used in conjunction with “an infringing copy”, i.e. “an article which is...an infringing copy”. This, say Getty Images, reinforces the wide meaning of “article”, since the word “copy” is sufficiently wide to cover both tangible and intangible things (as is borne out by sections 17(2) and 27(6) CDPA). Furthermore, they say that none of the activities which must be done with an article, as contemplated in sections 22 and 23 CDPA (e.g. importing, possessing, selling, distributing) are obviously aimed only at tangible things. These activities can also apply to an intangible article.
568. Stability, on the other hand, points to *Stroud’s Judicial Dictionary of Words and Phrases* which identifies various statutes and authorities in which the term “article” appears. Stability says that the term “article” is not used in any of these examples in a sense that would include intangible matter or abstract information. Contrary to Getty Images’ submissions, Stability says that the words used to define types of secondary dealing in sections 22 and 23 CDPA in fact point towards the word “article” meaning only a tangible thing.
569. It is clear from the extract from *Stroud’s Judicial Dictionary* relied on by Stability that, as one might expect, the word “article” will almost invariably take its colour from its surrounding context. On its own, it is not a word that has a clear meaning; it is no more than a generic term for an object, item or thing. It is commonly used in connection with tangible objects (such as an article of clothing or a household article) but that does not mean that it is incapable of applying to something intangible in the right context – for example a newspaper article or an article of faith. That the term is used in many statutes

or contracts to refer to a moveable chattel or physical object of some kind does not appear to me to assist in connection with this particular statute.

570. Sections 22 and 23 CDPA provide the colour. They apply to an “article” which is “an infringing copy of the work”. Such an “article” will be capable of being imported, possessed, sold, let for hire, offered or exposed for sale or hire, exhibited or distributed. However, it seems to me to be implicit in these sections that any article which falls within the definition of an “infringing copy” is intended for the purposes of this statute to be capable of being subject to these secondary types of dealing. Accordingly, to my mind, the first question for the court must be whether an infringing copy is capable of encompassing an intangible article. Arguments to the effect that an intangible article cannot be imported or possessed (although I shall return to them in a moment), do not really seem to me to assist in the abstract.

571. It is common ground that sections 22 and 23 CDPA are to be construed in their statutory context. As Stability submits, this includes the context provided by Part 1 to the Act which is concerned with copyright. Section 27 CDPA provides the meaning of infringing copy. It makes clear that the words “infringing copy” in Part 1 are to be construed in accordance with “this section” and I have set out sections 27(1)-(3) above. The section is concerned with the circumstances in which an article will amount to an infringing copy (whether because its making constituted an infringement of copyright or because it has been imported into the United Kingdom and its making in the United Kingdom would have constituted an infringement). The section does not deal with what is meant by “copying” (or indeed “copies” or “copy”) – for that one must look to section 17 CDPA.

572. Section 17 provides (insofar as relevant) as follows:

“17.— Infringement of copyright by copying.

(1) The copying of the work is an act restricted by the copyright in every description of copyright work; and references in this Part to copying and copies shall be construed as follows.

(2) Copying in relation to a literary, dramatic, musical or artistic work means reproducing the work in any material form. **This includes storing the work in any medium by electronic means.**

...

(6) Copying in relation to any description of work includes the making of copies which are transient or are incidental to some other use of the work” (**emphasis added**).

573. As I said in my judgment on the Summary Judgment application, it is clear from section 17(2) that reproduction of the copyright work (for the purposes of primary infringement) can be “in any material form” and that it includes “storing the work in any medium by electronic means”. Thus Getty Images contend that owing to the use

of Copyright Works in the training data (which had to be reproduced and stored locally or in the cloud), the making of Stable Diffusion involved reproducing the Copyright Works in a material form which (had it taken place in the United Kingdom) would have been contrary to section 17 CDPA 1988, because the unauthorised storage in electronic form of a copyright work will amount to an infringing reproduction of that work.

574. It seems to me that when construing the meaning of “infringing copy” and also, therefore, “article” for the purposes of sections 22, 23 and 27 CDPA, one cannot ignore the context provided by section 17 as to what may be an infringing copy. The words of the statute must be construed having regard to their proper context so as to ensure coherence and consistency.

575. As Getty Images submit, this approach is consistent with the approach adopted by Laddie J in *Sony v Ball*. In that case it was common ground that a Random Access Memory chip (or RAM chip) was capable of being an article, but the key question was whether a RAM chip that contained an infringing copy of digital data for only a fraction of a second could be an article (or as counsel put it, a “tangible substance”) given its ephemeral state. Laddie J held that it could. In setting out his reasoning, he referred to both sections 27 and 17 CDPA 1988 saying this:

“15. ...The silicon RAM chip is an article. When it contains the copy data, it is also an article. The fact that it did not contain the copy before and will not contain the copy later does not alter its physical characteristics while it does contain a copy. It is always an article but it is only an infringing article for a short time. There is nothing in the legislation which suggests that an object containing a copy of a copyright work, even if only ephemerally, is for that reason to be treated as not an article. On the contrary, the definition in s.27 points to the instant of making of the copy as crucial to the determination of whether or not it is an infringing article. An article becomes an infringing article because of the manner in which it is made. Whether it is an infringing article within the meaning of the legislation must be determined by reference to that moment. It matters not whether it is remains in that state, since retention as a copy is no part of the definition in the section

16. This is also consistent with the provisions of s.17...

17. In my view the meaning of the term “infringing copy” takes colour from its context in the Act. Just as a transient act of copying amounts to infringement, so an article which transiently contains a copy is an infringing copy for the purpose of this legislation. Taken together, these two subsections appear to suggest that even making a transient copy of a work can constitute making a reproduction “in a material form”. Thus RAM containing a copy of Sony’s copyright work is a reproduction in material form. **It would produce an unwarranted inconsistency in the Act were that material form not to be considered an article for the purpose of s.27**” (emphasis added).

576. This final sentence is strongly supportive of the proposition that the concept of an “article” as used in the statute cannot be divorced from the concept of an infringing copy, or indeed from what can amount to an act of copying or a copy under the Act.

577. Laddie J was not in that case directly concerned with (i) the question of what the word “article” means under the Act (there was no suggestion that a RAM chip could not be an article), or (ii) whether a copy stored on an intangible storage medium could still be an infringing copy (a RAM chip being a physical tangible means of storage). It does appear, however, that the argument proceeded on the assumption that the infringing copy must be tangible – hence the centrality of the question of whether a RAM chip containing a copy “is too short lived to be regarded as tangible”. Thus the judgment sheds some light on what amounts to an infringing copy when a work is stored by electronic means. Laddie J regarded the RAM chip as containing a copy of the work in “any material form” for the purposes of section 17(2). He found that the RAM chip (on which the work had been stored electronically, albeit for a short space of time), was an infringing copy, but it was only an infringing copy while the work was contained in it. Implicit in this judgment is an interpretation of infringing copy which (in the case of storage of an electronic work) includes (as Stability contends) the media on which it is stored.

578. Must that media inevitably be tangible? *Sony v Ball* is not authority for such a proposition. Laddie J did not directly address that question. The words “in any material form” in section 17(2) are addressed at the reproduction of the copy, while the concept of storage in “any medium” could hardly be more broad or generic.

579. At the time of the CDPA – as is common ground – storage media capable of storing electronic copies would have involved physical, tangible means of storage such as a floppy disk or magnetic tape (more recently a hard drive, flash drive or USB stick and, by the date of *Sony v Ball*, a RAM chip). However, in today’s world, electronic copies can be stored in an intangible “cloud” with no physical instantiation. Given the breadth of the words “any medium”, it seems to me that section 17(2) was plainly intended by Parliament to cover any means of storing electronic copies and thus must be capable of including modern means of storage such as cloud storage.

580. I agree with Getty Images that the “always speaking principle” is of assistance in these circumstances. Stability does not suggest that the statute was intended to be “frozen” in time and I consider that modern storage methods in intangible media amount to a fresh set of facts which fall within the same genus of facts as those to which the original expressed policy has been formulated. The fresh set of facts arises by reason of the prevalence in the modern world of intangible electronic storage which has been brought about by enormous strides in technology since the date of commencement of the CDPA. The purpose of the Act – the protection of copyright owners – would, in my judgment, be fulfilled by an interpretation which encompassed modern technology.

581. Stability’s submission that electronic copies of copyright works were known to Parliament in 1988 is true, but tilts at the wrong windmill. On Stability’s own submissions, intangible means of storage had not yet been invented, so the failure to make express provision for them in the Act cannot possibly be said to be a deliberate policy. I reject Stability’s suggestion that I am in any way assisted on this question by the provisions in the CDPA dealing with broadcasting and transmission, which have nothing whatever to do with the storage of electronic copies.

582. I am fortified in this view by reference to the fact that, in a slightly different context, the CJEU has held that saving a copy of a work in storage space made available to a user in connection with cloud computing services constitutes a reproduction of that work (see *Austro-Mechana Gesellschaft zur Wahrnehmung mechanisch-musikalischer Urheberrechte GmbH v Strato AG* (C-433/20) EU:C:2022:217, [2022] ECDR 10, a case on Article 5(2)(b) of Directive 2001/29, at [15]-[18]).

583. In the circumstances, an electronic copy stored in an intangible medium (such as the AWS Cloud) is, in my judgment, capable of being an infringing copy and thus also capable of being “an article”. I cannot see that Parliament could objectively have intended the use of the generic term “article” in sections 22, 23 and 27 CDPA to limit the scope for protection in relation to electronic copies (as provided for in section 17), including in circumstances where the means of storing such copies might change by reason of advances in technology.

584. I must address a few of Stability’s additional arguments. While on the one hand accepting that the words of a statute must be seen in their proper context, Stability argues that the provisions of section 17 are of no assistance in construing the term “infringing copy”, because “an infringing copy” must be not just a copy but also an article under each of sections 22, 23 and 27 CDPA”. Stability submits that the requirement that the “copy” is an “article” is “an additional requirement which means that the secondary infringement provisions can only apply to copies in tangible form”. This is effectively just a different way of putting the point I have addressed above about the storage medium for electronic copies forming part of the infringing copy. It takes matters no further. In addition, to my mind this submission looks at the statute from the wrong end of the telescope. The relevant sections refer to “an article which is...an infringing copy” – this drafting neither suggests nor implies that in addition to being a copy, an infringing copy must also be something else – on Stability’s case, an article in tangible form. In my judgment this interpretation is neither consistent with the words of sections 22, 23 and 27, nor with a proper understanding of those words in their context (which includes section 17).

585. I am also unpersuaded by Stability’s argument that the various acts of secondary infringement can only apply where there is a tangible article. While the concepts of “importation” and “possession” (on which Stability focused) may ordinarily connote physical transfer or control of a tangible object, I am of course only concerned with those concepts as they are used in the context of this statute. None of the cases to which I was referred by Stability assists on that point:

- i) As to “importation”, Stability relied on *LA Gear Inc v Hi-Tec Sports Plc* [1992] FSR 121, in which Morritt J observed at [30] that “the importation must have occurred when the goods were physically received in this country...”, but a reading of the judgment confirms that this observation is expressly tied to the facts of that case (which concerned the importation of shoes). Morritt J was not concerned with the scope of the concept of importation for the purposes of sections 22 and 23 CDPA.
- ii) As to “possession”, Stability relied upon authorities concerned with common law liens or claims in detinue and conversion for the propositions that intangible information cannot be “possessed” (see e.g. *Your Response Limited v Datateam Business Media Limited* [2014] EWCA Civ 281, [2015] FSR 23 (“Your

Response”)) and that there is no property in information (see *Thaler v Comptroller General of Patents, Trade Marks and Designs* [2021] EWCA Civ 1374, [2021] RPC 19 per Arnold LJ at [125]). However, while *Your Response* certainly makes plain that (consistent with the decision of the House of Lords in *OBG Ltd v Allan* [2007] UKHL 21) “the common law draws a sharp distinction between tangible and intangible property” (per Moore-Bick LJ at [15]), the Court of Appeal in that case was not considering the meaning of “possession” in the context of the CDPA. Indeed, in a paragraph on which Stability partially relied (at [42]), Floyd LJ expressly acknowledged the scope for intangible information to give rise to intellectual property rights, such as database right and copyright, notwithstanding that more generally the law has been reluctant to treat information itself as property.

586. Further, that an “article” which is an infringing copy can be “intangible” in that it can encompass electronic copies (which may not be “stored” in any tangible form), finds support, as Getty Images submits, in section 27(6) CDPA which provides that:

“In this Part “infringing copy” includes a copy falling to be treated as an infringing copy by virtue of any of the following provisions”

By the time of closing submissions, it was accepted by Getty Images that only one of the provisions listed under s. 27(6) was relevant as having been present in the statute at the time of its enactment (see *R (Brown) v Secretary of State for the Home Department* [2015] UKSC 8, [2015] 1 WLR 1060 per Lord Toulson JSC at [24]), namely:

“Section 56(2) (further copies, adaptations, &c. of work in electronic form retained on transfer of principal copy),...”

587. Section 56 (insofar as relevant) provides as follows:

“56.— Transfers of copies of works in electronic form.

(1) This section applies where a copy of a work in electronic form has been purchased on terms which, expressly or impliedly or by virtue of any rule of law, allow the purchaser to copy the work, or to adapt it or make copies of an adaptation, in connection with his use of it.

(2) If there are no express terms—

(a) prohibiting the transfer of the copy by the purchaser, imposing obligations which continue after a transfer, prohibiting the assignment of any licence or terminating any licence on a transfer, or

(b) providing for the terms on which a transferee may do the things which the purchaser was permitted to do,

anything which the purchaser was allowed to do may also be done without infringement of copyright by a transferee; but any

copy, adaptation or copy of an adaptation made by the purchaser which is not also transferred shall be treated as an infringing copy for all purposes after the transfer”.

588. As Stability argues, section 56(2) is a deeming provision: it deals with the situation where a purchaser (who has obtained a copy of an electronic work which he can legitimately copy for his own purposes pursuant to the terms of purchase) transfers the copy to a third party in the absence of any contractual provisions determining the then status of any copy or adaptations of the copy in the hands of the transferee or the original purchaser. Specifically it deems that any copy or adaptation made by the purchaser which is not transferred to the third party “shall be treated as an infringing copy” .
589. Accordingly, in my judgment, Stability is right that this provision is designed to address a scenario that might “fall through the cracks” in the sense that it would otherwise be unclear what rights the transferee might have in the copy of the work upon transfer and what the status of any copy of the work retained in the hands of the purchaser following transfer might be. However, I disagree with Stability that section 56(2) is supportive of the proposition that an electronic copy could not be an infringing copy absent this deeming provision. To my mind, the section is not designed to address the question of whether an electronic copy can be an infringing copy (that is already plain from section 17). All this section is doing is ensuring that, in the specific factual scenario postulated, an electronic copy shall be treated as an infringing copy, notwithstanding that the purchaser originally obtained the right to copy the work. Interestingly, Stability submits that any copy retained by the purchaser “would [in 1988] inevitably involve storage on a physical medium” – but that of course circles back into the question I have already addressed.
590. Pausing there to take stock, on a proper construction of the statute, I consider that an article, which must be an infringing copy, is capable of being an electronic copy stored in intangible form. Standing back, I agree with Getty Images that if the word “article” were construed as only covering tangible articles, this would deprive authors of protection in circumstances where the copy is itself electronic and it is then dealt with electronically. Not only would that be inconsistent with the words of the statute, but it would also be inconsistent with the general scheme of copyright protection which is to reward authors for their creative efforts. As Lord Bingham said in *Designers Guild Ltd v Russell Williams (Textiles) Ltd* [2000] 1 WLR 2416 at 2418:

“The law of copyright rests on a very clear principle: that anyone who by his or her own skill and labour creates an original work of whatever character shall, for a limited period, enjoy an exclusive right to copy that work. No one else may for a season reap what the copyright owner has sown”.

As Getty Images submit, the importance of this principle is particularly acute in the modern world where huge numbers of works and copies are in electronic, intangible, form, owing to developments in technology. The “always speaking” principle of interpretation appears to me to support a construction that facilitates protection for copyright owners.

591. I reject Stability’s contention that the relevant legislative provisions (insofar as they concern the interpretation of “article”) are “ambiguous, obscure or, on a conventional

interpretation, lead to absurdity” such that pre-legislative statements made in Parliamentary debates are admissible on this question of statutory construction (see *R (Coughlan) v Minister for the Cabinet Office* [2022] UKSC 11, [2022] 1 WLR 2389 (“*Coughlan*”) per Lord Stephens JSC at [14] referring to the well-known conditions set out by Lord Browne-Wilkinson in *Pepper v Hart* [1993] AC 593 at 640). While the interpretation of the word “article” is not entirely straightforward and, certainly in this case, has led to considerable debate, I do not consider that a close analysis of the statute (the primary source for interpretation) and an application of the well-known principles of statutory construction to which I was referred, gives rise to any ambiguity or obscurity. Accordingly, there is no need for me to extend an already very lengthy judgment by considering the extracts from Hansard to which Stability drew my attention and I reject Stability’s submission that some form of “reverse engineering” in relation to Hansard (involving looking at Hansard to see whether it is inconsistent with Getty Images’ submissions before determining whether there is any ambiguity in the statute) would be appropriate. Even if I am wrong about that, I am inclined to agree with Getty Images that the Hansard materials on which Stability seeks to rely are not “clear and unequivocal on the point of interpretation which the court is considering” (*Coughlan* at [14]) owing to their apparent focus on broadcasts and cable programmes.

The Interpretation of “an infringing copy”

592. That an “infringing copy” may be an electronic copy stored in an intangible medium, such that an “article” for the purposes of sections 22 and 23 CDPA encompasses an electronic, or intangible copy, is not however the end of the argument. Stability contends that whether an article is tangible or intangible, an infringing copy must be a copy; i.e. a reproduction of a Copyright Work. This is central to Stability’s case because it says that where, as is common ground, the model weights of the various Stable Diffusion versions do not store the visual information in the Copyright Works, they cannot amount to a copy.
593. Getty Images say that Stability’s interpretation is obviously wrong, because the words of section 27(2) and (3) CDPA make plain that what is required if an article is to be an infringing copy is that “its making constituted an infringement”. They submit that there is no requirement that the article thus made must continue to retain a copy or copies of the work and they point to the fact that section 17(6) makes clear that “copying in relation to any description of work includes the making of copies which are transient or are incidental to some other use of the work”. They continue to emphasise the importance of a purposive construction. They also rely upon *Sony v Ball* at [15] (set out in full above) focussing specifically on Laddie J’s observations that:

“[a]n article becomes an infringing article **because of the manner in which it is made**. Whether it is an infringing article within the meaning of the legislation **must be determined by reference to that moment. It matters not whether it remains in that state, since retention as a copy is no part of the definition in the section**” (emphasis added).

594. At the heart of Getty Images’ case is the submission that, unlike the facts of *Sony v Ball*, “[i]n this case the time of making of the copies of the Copyright Works **coincides** with the making of the article, i.e. the AI model. Thus unlike the [RAM] chip, **as soon as it is made**, the AI model is an “infringing copy”” (emphasis added).

595. Sections 22 and 23 CDPA are not addressed to the act of reproduction of a copy, but to the downstream dealings in an article which is “an infringing copy”. Section 27 provides for the meaning of “infringing copy” which is to be construed “in accordance with this section”. This led Getty Images in opening to submit that the scope for looking at the statutory context (at least for the purposes of construing the phrase “infringing copy”) is limited. I disagree, as will already be clear. Stability points out (in my view, correctly) that “infringing copy” is a composite phrase and that the fact that it is defined as a phrase does not negate the meaning of the constituent words. It drew my attention to *G4S Plc v G4S Trustees Ltd* [2018] EWHC 1749 (Ch) in which Nugee J (as he then was) had to construe the term “pensionable service”, which was a defined phrase in the Pensions Act 1995. At [26] he said this:

“I am conscious that in section 124(1) of the Pensions Act 1995 the expression “pensionable service” is itself a defined expression, and that ultimately what I have to construe are the words of the definition, not the words “pensionable service” as such. But where a statute (or indeed a contract) uses a defined expression, the ordinary meaning of the word or words defined is part of the material which can be used to construe the definition. The defined expression is seldom an arbitrary label; it is usually chosen as a distillation of the meaning of a concept more precisely stated in the definition: see *Birmingham City Council v Walker* [2007] 2 AC 262, para 11, per Lord Hoffmann and compare *Chartbrook Ltd v Persimmon Homes Ltd* [2009] AC 1101, para 17, again per Lord Hoffmann.”

596. I agree with Stability that Nugee J’s reasoning is of direct application here. As I have already indicated, the word “copy” is not arbitrary – the statute provides in section 17 (entitled “Infringement of copyright by copying”) that it shall have a particular meaning and it seems to me that that meaning must inform the proper interpretation of the composite phrase “infringing copy”. This is entirely consistent with the approach of Laddie J in *Sony v Ball* as set out above, who expressly considered (at [15]) that “the meaning of the term “infringing copy” takes colour from its context in the Act”. It is also consistent, as Stability points out, with the fact that the concept of copying and copies is a fundamental concept in an Act which is concerned with the provision of protection against the copying of copyright works.

597. Thus, it seems to me to be clear that an infringing copy must be a copy, as Stability submits; the essence of the infringement is that there has been an infringement of copyright by the reproduction of the work (including by its storage in any medium by electronic means) in any material form. Consistent with the decision in *Sony v Ball*, I cannot see how an article can be an infringing copy if it has never consisted of/stored/contained a copy. In *Sony v Ball* the RAM chip (which was always an article) was only an infringing copy while it contained the copy but not otherwise: it was therefore, in the words of the Judge, “only an infringing article for a short time”, i.e. for the time that it contained the copy.

598. As *Sony v Ball* also makes plain, an 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. Thus it was no part of Stability’s case, (notwithstanding that Getty Images wrongly sought to characterise

Stability's case in this way), that for an article to be an infringing copy it must continue to retain a copy or copies of a Copyright Work. Stability accepts that transient copying can be sufficient to give rise to an infringing copy, but it says that an infringing copy must (at least at some point) contain a copy.

599. In reality therefore, the dispute between the parties as it finally emerged in closing, really turns on whether an article whose making involves the use of infringing copies, but which never contains or stores those copies, is itself an infringing copy such that its making in the UK would have constituted an infringement. Taking the specific facts with which I am concerned, 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?
600. In my judgment, it is not. It is not enough, as it seems to me, that (in Getty Images' words) "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. They are purely the product of the patterns and features which they have learnt over time during the training process. Getty Images' central submission that "as soon as it is made, the AI model is an infringing copy" is, accordingly, in my judgment, entirely misconceived. Unlike the RAM chip in *Sony v Ball* which became an infringing copy for a short time, in its final iteration Stable Diffusion does not store or reproduce any Copyright Works and nor has it ever done so. The fact that its development involved the reproduction of Copyright Works (through storing the images locally and in cloud computing resources and then exposing the model weights to those images) is of no relevance. Furthermore, that there is no requirement that an article which is an infringing copy must continue to retain a copy does not assist Getty Images, because it is implicit in the word "continue" that at some point the article has in fact contained an infringing copy. The model weights for each version of Stable Diffusion in their final iteration have never contained or stored an infringing copy.
601. I agree with Stability that the concept of an infringing copy cannot be interpreted in the abstract without reference to the fundamental nature of a copy, as revealed by section 17. Such an interpretation precludes the potential for an article which *has never contained a copy* to be capable of being "an infringing copy". But, sections 27(2) and (3) require the article that is made to be a "copy". They are not concerned with a process which (while it may involve acts of infringement) ultimately produces an article which is not itself an infringing copy. Section 27(6) provides no assistance to Getty Images in this regard. I agree with Stability that when applying the 'making' concept to an intangible article¹⁰, the conception or creation of the model weights (i.e. the process of creation of information) cannot be the relevant act of making – that can only be the act of production of the resulting intangible record of that information, which derives from that process of creation. It is that intangible record which, on Getty Images' case, enters

¹⁰ I accept that this is more difficult than would be the case with a tangible article, albeit that does not seem to me to be a reason for construing an article as involving only tangible objects, just as it is not a reason to treat the concept of "making" in the broad terms suggested by Getty Images.

the UK and that is the specific article whose making (had it taken place in the UK) must have constituted an infringement of copyright.

602. Finally, on this issue I am not persuaded that Getty Images' "purposive interpretation" assists. The scheme of the statute is to provide protection to the owners of copyright works against acts of infringement involving copying. Ms Lane highlighted the importance of being able to prevent people from circumventing copyright protection by carrying out infringing acts abroad and then importing the fruits of those acts into the United Kingdom. However, where an AI model itself is not an infringing copy, it cannot have been the intention of Parliament that it should fall within the meaning of infringing copy in section 27. The secondary acts to which Getty Images objects thus fall outside the policy and object of the Act.

The factual questions

603. In light of my findings on the law, it is not strictly necessary for me to make any determinations on the facts. I have already found that Stable Diffusion is not capable of being an infringing copy and thus sections 22 and 23 CDPA are not applicable. However, in case I am wrong in my interpretation of the CDPA, I nevertheless set out my findings on the facts as briefly as possible in this section.

604. On the assumption that, contrary to my findings above, Stable Diffusion is capable of being an infringing copy for the purposes of section 27(3), then:

- i) There is evidence that Copyright Works were used for training the Model and Stability has made limited concessions to the effect that "at least some" Copyright Works were used in training the Model and that images from the Getty Images Websites were used in training. In a Response to an RFI, Stability made clear that it would not seek to challenge at trial the proposition that sample works A1-A17 from the SOCI were used in training versions 1.4, v2.x, all versions of XL and v1.6.
- ii) I find that, as Stability acknowledged in closing, the model weights for Stable Diffusion v2.x, XL 0.9, XL Base 1.0 and XL Turbo have all been made available by Stability for download on Hugging Face. This involves making intangible copies of these Models available for download onto a local device in the UK. Given that I have found that an article may be intangible, I accept that downloads of the Model amounted to importation for the purposes of the Act and that they have taken place within the United Kingdom. Stable Diffusion is an open source model which is available for anyone to download. Mr Auerhahn agreed in cross-examination (and I find) that there will have been many downloads of Stable Diffusion v2.x and SD XL on to computers in the UK (perhaps running into millions). This was not disputed by Stability in closing submissions. Under section 27(3) the importation need not be by the defendant and so it matters not whether the download is characterised as an act of the defendant or otherwise.
- iii) Getty Images accept (and I find) that gaining access to the Model via DreamStudio does not involve importation or even transfer of a copy of Stable Diffusion to the UK because DreamStudio is a hosted service that is provided from servers located outside the UK. No copy of the Model is ever provided to the user and all inference and output synthesis takes place outside the UK.

Whatever my decision on the interpretation of the statute, there could never be any act of secondary infringement by reason of the provision of remote software services. The case on secondary infringement under sections 22 and 23 CDPA must fail in relation to DreamStudio. Given the way in which the Developer Platform works, the same must equally apply to it. In closing there was some discussion about Stability's case that, since 14 December 2023, SD XL Turbo has been made available for download on new licensing terms. This appears to have been part of an AI model licensing programme which (given the period it has been available) I find on balance has involved downloads in the UK.

iv) As I understood Getty Images' submissions on the facts in closing, it contends that it is "clear that, at the very least, Stability has distributed Stable Diffusion in the course of business via Hugging Face". I agree in relation to v2.x, SD XL and v1.6. I accept that Stability is responsible for these acts of distribution which (if the other statutory requirements are met) thereby amount to acts of secondary infringement pursuant to section 23(d) CDPA. The same applies in relation to downloads from the Model licensing programme.

605. As to the question of whether (if Stable Diffusion is an infringing copy) Stability had knowledge or reason to believe that it was, Stability's Defence focuses on the fact that Stability did not know or believe that the Model reproduced any element of the training dataset or any of the images contained within it and that it did not otherwise know of any facts or circumstances which would have put a reasonable person on notice of the contrary. However, I can only address this issue on the assumption that Stable Diffusion is an infringing copy because of the training process (i.e. that, contrary to my finding above, the reproduction and use of infringing copies during training is sufficient to satisfy the statutory test under section 27(3)). On that basis, the point can be dealt with shortly by simply setting out what Stability did know or have reason to believe.

606. Although Stability failed to call any evidence from a witness designed to address the issue of knowledge, in my judgment there is available evidence that Stability knew or had reason to believe:

- i) that works of the sort with which this case is concerned are protected by copyright and that acts of reproduction, communication to the public, importation and possession of copies in the course of business are prohibited without the consent of the copyright owner – this much was admitted by Stability in its Defence. Mr Konstantinov confirmed that "copyright was something that people were talking about generally in the AI world", while Ms Bakshandaeva confirmed in her evidence that she thought she was aware that stock images were protected by copyright;
- ii) that Stable Diffusion was being trained on LAION datasets and that these had been produced by filtering the Common Crawl public web archive for images and storing them alongside their HTML alt-text (as confirmed by the Model Cards, by an internal Finance Department document dated 27 June 2023 and by the evidence of Professor Farid);
- iii) that the training data had been scraped from the internet without the consent of copyright owners; and

iv) that the training data included works bearing the Getty Images watermark – had that not been the case, Stable Diffusion could not have produced output bearing the watermarks*. Mr Wagrez confirmed that he understood that it was “problematic having Getty Images in the dataset.”

607. Furthermore, I accept Getty Images’ submissions that the latter two points are supported by disclosure documents from the trial bundles. In addition to evidence to which I have already referred in this judgment, I refer specifically to:

- i) an internal Stability Chat exchange taking place on 7 June 2022 in which the individuals involved in the chat (including Mr Mostaque) acknowledge the need to get “watermarks out of the dataset”, the potential (depending on “the PR/legal fight”) for it to be possible only to train on “public domain data”, the fact that “...people will go ‘this model was trained on our copyrighted data and we were not compensated for its use’ when they perceive it to threaten their livelihoods” and that “whether or not we turn out to be legal in the clear for using it in the end, there is going to be a ton of noise over it”.
- ii) An internal Stability Chat exchange taking place on 4 February 2023 in which the individuals involved in the Chat (including Mr Rombach) discuss options for “watermark removal” from the training set and refer to images with watermarks as making up “half the data”.
- iii) An email of 8 February 2023 from Joe Penna of Stability to Kenneth Lee and others at Stability, stating “At the moment we’re finding every model that Stability has available in staging & development (DeepFloyd, Stable Diffusion XL...) is not releasable due to issues with: Copyright (replicating logos, copyrighted artworks, etc etc etc...)”,
- iv) An internal Stability Chat exchange on 10 March 2023 in which Conner Ruhl of Stability refers to experiencing “watermarks and furniture...20% of the time”.
- v) An internal Stability Chat exchange from 6 April 2024 in which Mr Vencu explains that “urls with getty in their domain are clearly coming from a getty website”.

608. Furthermore, a letter of claim was sent to Stability on 16 January 2023 asserting that Getty Images Content (said to be 12 million links) had been scraped from the internet and used to train AI models without consent which were now available to users via DreamStudio, thereby infringing copyright (amongst other things). The Particulars of Claim were first served on 12 May 2023.

609. In closing, Ms Lane invited me to draw an inference in relation to the issue of knowledge by reason of Stability’s decision not to call Mr Vencu. While this is an issue in respect of which there might have been more scope for such an inference, I do not consider it would have been necessary, given the evidence to which I have referred above. Furthermore, I note that the Agreed Decision Tree for the Secondary Infringement Claim provided by the parties after the trial indicates that knowledge is admitted by Stability in the event that Getty Images’ interpretation of the CDPA were

to be accepted. I had not understood that to be the position at trial, which is why I have made the necessary findings of fact.

(I) COPYRIGHT SUBSISTENCE AND OWNERSHIP

610. Notwithstanding that my decision that Stable Diffusion is not an infringing copy for the purposes of the CDPA means that the claim of secondary infringement falls to be dismissed, I must nevertheless address the outstanding issues which might ultimately prove important if I am wrong in my interpretation of the statute. First, I must consider the limited Ownership Issues that remain. Issues of subsistence and ownership for the purposes of the Secondary Infringement Claim are to be decided on the basis of a sample, namely the SOCI Works and Sample Works A-K.

611. By the end of trial:

- i) Getty Images had abandoned its claims in relation to SOCI Works A1, A12, A13 and A17 and so I need not address those further.
- ii) there was also no issue between the parties on authorship or dates of creation of any of the remaining SOCI Works or Sample Works A-K and no issue as to subsistence of copyright in these works, which is admitted by Stability.
- iii) there was no challenge to Getty Images' title to copyright in SOCI Works A2, A5, A6, A7, A8 and Sample Works A-D.
- iv) there was no challenge to the licensor's title to copyright in SOCI Works A14-A16 and no challenge to the Sixth Claimant's title to copyright in Sample Works E-K.

612. Although this was sufficient for the purposes of Getty Images' Secondary Infringement Claim, Ms Bowhill invited me nevertheless to make findings in relation to the limited number of remaining works that remained in issue between the parties, essentially to facilitate submissions that Getty Images may wish to make in relation to relief in due course. Although Stability rejects any suggestion that the SOCI Works and the Sample Works A-K form a genuine "sample" or that there is any basis in the directions of the Court or the statements of case for any extrapolation to be made from the works in issue, I did not understand it to object to my making findings on the remaining works. As they have been fully argued, it makes good sense for the court to determine the factual position in relation to each.

613. The sample works in respect of which title remains in issue are SOCI Works A3, A4 and A9-11. Getty Images allege that they own the copyright in each of these works.

Legal Framework

614. The relevant law was uncontroversial and I gratefully adopt a summary of the law used by Getty Images in its opening skeleton argument.

615. Section 1(1)(a) of the CDPA provides, so far as material, as follows:

“(1) Copyright is a property right which subsists in accordance with this Part in the following descriptions of work—

- (a) original literary, dramatic, musical or artistic works...
- (b) sound recordings, films or broadcasts...”

616. Artistic works are defined in section 4(1)(a) of the CDPA as follows:

“(1) In this Part “artistic work” means—

- (a) a graphic work, photograph, sculpture or collage, irrespective of artistic quality,
- (b) a work of architecture being a building or a model for a building, or
- (c) a work of artistic craftsmanship.”

617. Films are defined in section 5B(1) as “a recording on any medium from which a moving image may by any means be produced”.

618. Authorship is dealt with in section 9 of the CDPA, which provides in sub section (1) that “author” in relation to a work means “*the person who creates it*”.

619. In the context of photographs, there is generally no difficulty in identifying the author. On the facts of the present case there is no dispute that it will be the “photographer”.

620. Section 11 of the CDPA provides for first ownership of copyright as follows:

“(2) The author of a work is the first owner of any copyright in it, subject to the following provisions.

(3) Where a literary, dramatic, musical or artistic work or a film, is made by an employee in the course of his employment, his employer is the first owner of any copyright in the work subject to any agreement to the contrary.”

621. The terms “employee”, “employer” and “employment” are defined in s 178 of the CDPA as referring to “*employment under a contract of service or of apprenticeship*”.

622. The concepts of an employee and a contract of service have the same meanings in the CDPA as they do elsewhere: *Ultraframe (UK) Ltd v Fielding* [2003] EWCA Civ 1805, [2004] ECDR 34 at [19].

623. The three conditions necessary for a contract of service or contract of employment to exist were confirmed by the Supreme Court in *Autoclenz Ltd v Belcher* [2011] UKSC 41, [2011] ICR 1157 at [18]-[19]¹¹ (citing *Ready Mixed Concrete (South East) Ltd v Minister of Pensions and National Insurance* [1968] 2 QB 497):

¹¹ Affirmed by the Supreme Court in *HMRC v Professional Game* [2024] UKSC 29 at [27]-[28].

“18 ...A contract of service exists if these three conditions are fulfilled. (i) The servant agrees that, in consideration of a wage or other remuneration, he will provide his own work and skill in the performance of some service for his master. (ii) He agrees, expressly or impliedly, that in the performance of that service he will be subject to the other’s control in a sufficient degree to make that other master. (iii) The other provisions of the contract are consistent with its being a contract of service . . . Freedom to do a job either by one’s own hands or by another’s is inconsistent with a contract of service, though a limited or occasional power of delegation may not be...

19 Three further propositions are not I think contentious: (i) ...There must . . . be an irreducible minimum of obligation on each side to create a contract of service. (ii) If a genuine right of substitution exists, this negates an obligation to perform work personally and is inconsistent with employee status... (iii) If a contractual right, as for example a right to substitute, exists, it does not matter that it is not used. It does not follow from the fact that a term is not enforced that such a term is not part of the agreement...”

624. Finally, where a copyright work is created by a director of a company who is not employed under a contract of service, it may result in the director holding the copyright work on trust for the company such that the copyright must be assigned to the company if the director is called upon to do so: see *Copinger* at [4-214]. As cited by Mr David Stone in *Mei Fields Designs Ltd v Saffron Cards and Gifts Ltd* [2018] EWHC 1332 (IPEC), [2018] FSR 33 at [65] (referring to the judgment of Richard Arnold QC (as he then was) in *Vitof Ltd v Alloft* [2006] EWHC 1678 (Ch) at [144]):

“There is, however, no rule that works created by a director for his company are always held on trust: it will depend on what, if anything has been agreed. In particular, it is always open to the shareholders of a company to agree that a director should retain property he has created or to relieve him of any liability for any breach of duty, provided that to do so is not ultra vires the company or a fraud on its creditors.”

625. I now turn to deal with each of the disputed works in turn.

The Disputed Works

Works A3 and A4

626. These works (both images of storm clouds gathering at sunset at the Gabba cricket ground), referred to earlier as the Gabba Images, were created by Mr Kolbe on 7 November 2015. Getty Images contend that, at the time the Gabba Images were taken, Mr Kolbe was an employee of Getty Images Sales Australia Pty Ltd (“**Getty Sales Australia**”), which was therefore the first copyright owner. They say that copyright

was then assigned to the Second Claimant pursuant to an Imagery Development and Marketing Services Agreement (“**the IDMSA**”) dated January 2016 (in respect of which there was a Confirmatory Assignment between Getty Sales Australia and the Second Claimant on 17 September 2024) and that by a Merger Agreement dated 24 June 2016 (“**the Merger Agreement**”) copyright in the Gabba Images was assigned by the Second Claimant to the First Claimant.

627. I need not go into the detail of these transactions because the only issue between the parties is whether Mr Kolbe (who was not called to give evidence) was in fact employed by Getty Sales Australia at the time the Gabba Images were created. Stability contends that the documents show that he was not, and I agree.
628. The only available contract of employment (dated 17 June 2002) is between Mr Kolbe and a company called Getty Images Pty Limited (“**GIPL**”). The contract confirms that from November 1998, Mr Kolbe was employed as a picture editor for the company working from its offices. It also confirms that any intellectual property created by Mr Kolbe in the course of his employment “is assigned to [GIPL]”. Mr Kolbe signed this contract on 5 August 2002.
629. The evidence confirms that GIPL and Getty Sales Australia have different company numbers and different parent companies. Getty Sales Australia was not incorporated until 7 June 2006, eight years after the commencement of Mr Kolbe’s employment with GIPL. There is no contract or other contemporaneous documentation evidencing Mr Kolbe’s employment, at any time, by Getty Sales Australia.
630. Despite the absence of any contract between Mr Kolbe and Getty Sales Australia, Ms Cameron gave evidence in her statement that in December 2024 a Human Resources Coordinator at Getty Images had accessed a repository called Workday, i.e. a human resources information system. She said that he had confirmed that Workday recorded that Mr Kolbe had been hired on 2 November 1998, that his employer was Getty Sales Australia and that his “Termination Date” was 31 May 2024. This same information was set out in an “Employment Confirmation” letter by Briar Hawkins of Getty Sales Australia: “This letter confirms the employment of Mark Kolbe...at Getty Images Sales Australia Pty Ltd, between 2 November 1998 and 31 May 2024...”.
631. The difficulty with this evidence is that the Workday records suggest, and the Employment Confirmation Letter says in terms, that Mr Kolbe was employed continuously from November 1998 to May 2024 by Getty Sales Australia. This cannot possibly be the case owing to the date of incorporation of that company. Ms Cameron explained in her oral evidence that Getty Sales Australia was Mr Kolbe’s employer on the Termination Date and that “it is possible that there were entities which changed over time”.
632. To my mind this evidence does not begin to establish that Getty Sales Australia was, on balance, Mr Kolbe’s employer when the Gabba Images were taken. At best the information in the Employment Confirmation letter and the Workday records is incomplete; at worst, it is wrong. Ms Cameron’s evidence goes no further than possibilities.
633. Ms Bowhill drew my attention to the IDMSA of January 2016 which provides in recital E that Getty Sales Australia “has the operational structure, facilities, expertise and

experience to perform...development and marketing activities" on behalf of the Second Claimant. Getty Images assert that "[i]t follows that, as at January 2016 Getty Sales Australia was the entity operating the business of developing images in Australia which **must have included employing the relevant photographers who create those images**" (**emphasis added**). They point out that Recital E to the IDMSA terminates an earlier IDMSA dated 1 August 2006 between the same entities and say that it follows from this that as at August 2006 Getty Sales Australia was operating the business of developing and marketing images in Australia "including employing the photographers who created those images".

634. To my mind, however, these submissions are a step too far. There are no documents to support this theory and in my judgment there is no proper basis on which I could find that Mr Kolbe was, on balance, employed by Getty Sales Australia on the date he took the Gabba Images. The only contemporaneous evidence of his employment by a company in the Getty Images Group is to be found in the contract of employment dated June 2002. There is no evidence to establish if and when Mr Kolbe ever took up employment with a different entity and no evidence of any assignment by GIPL to one of the Claimants in these proceedings.
635. Accordingly I find that Getty Images are unable to make out title to the copyright in the Gabba Images, i.e. SOCI Works A3 and A4.

Work A9

636. This work (an image of a CTM interview between Glen Campbell and Kim Woollen) was created by Mr Rick Diamond (who has not been called to give evidence) on 19 September 2011. Getty Images contend that, at the time, Mr Diamond was employed by the First Claimant and was based in Los Angeles. Stability challenges this contention having regard to the available documents.
637. The only document in evidence concerning Mr Diamond is a Confidentiality, Non-Disclosure and Assignment of Inventions Agreement ("**the NDA**") dated 2 July 2007. The counterparty to this agreement, which is expressly said to be an employment "at-will" and not a contract of employment, is Getty Images, Inc, a Delaware Corporation. The available evidence indicates that Getty Images, Inc is the parent company of both the First and Second Claimant.
638. The Workday records referred to in Ms Cameron's statement apparently record that Mr Diamond was hired on 26 April 2007 by the First Claimant and that his Termination Date was 3 January 2018, but there is no contemporaneous contract of employment with the First Claimant and Ms Cameron was not aware of any other employment relationship. It is not clear how or why Mr Diamond could have entered into the NDA in July 2007 if, at the same time, he was already an employee of the First Claimant as Ms Cameron's evidence suggests, and no explanation has been provided by Getty Images.
639. However, screenshots from the Getty Images website record that Mr Diamond is a "Staff" creator of Work A9 and Ms Cameron explained in her evidence that Mr Diamond was associated with the Los Angeles office and that:

“I have knowledge in my work at Getty Images that photographers based in that office in Los Angeles were employed by [the First Claimant]. That is who sent them their pay cheque. I have knowledge of that from my work at Getty Images over many years”.

640. While the documentary evidence is inconclusive and while the circumstances surrounding the NDA are puzzling, I accept Ms Cameron’s oral evidence and I find in light of that evidence that, on balance, Mr Diamond was an employee of the First Claimant when Work A9 was produced. I did not understand Stability to raise any other argument by way of impediment to this conclusion.
641. Accordingly I find that Getty Images have established title to the copyright in SOCI Work A9.

Works A10 and A11

642. Works A10 (an image of Eric Dane) and A11 (an image of Christopher Nolan) were created by Mr Rentz on 21 May 2012 and 13 May 2018 respectively. At the relevant time, Getty Images allege that Mr Rentz was an employee of Getty Images Devco Deutschland GmbH (“**Getty Devco Deutschland**”), which was accordingly the first owner of the copyright in works A10 and A11. This copyright is alleged to have been assigned to the Second Claimant pursuant to the IDMSA dated January 2016 (and also pursuant to a Confirmatory Assignment dated 18 September 2024). Pursuant to the Merger Agreement, Getty Images contend that copyright was assigned from the Second Claimant to the First Claimant on 24 June 2016.
643. Once again I need not look at each of these transactions in detail, the issue between the parties appears to come down to the interpretation of the terms on which Mr Rentz assigned copyright to his employer.
644. Mr Rentz’s contract of employment dated 13 March 2005 is with a company called Getty Images Deutschland GmbH (“**Getty Deutschland**”) but a confirmation letter dated 29 June 2006 (i.e. long before the A10 and A11 images were taken) confirms the transfer of his employment to Getty Devco Deutschland. In his statement, Mr Rentz explains, and I accept, that he was given notice of this transfer on 25 May 2006 and that since 1 August 2006 his employer has been the First Claimant.
645. In its translated form, his contract provides at clause 9(1) that:

“Mr Rentz will put all his energy at the service of the company and will work exclusively for the Getty Images Group. The Company has the sole right to exploit, or have exploited, all work results (images) from the activity without any material, temporal or spatial restrictions”.

646. Stability contends that this is not an assignment of copyright but a licence. It submits that any other position would be inconsistent with the reference to “without any material, temporal or spatial restrictions” – if his employer was the owner of this copyright by assignment then “it would go without saying that its rights were “without any material, temporal or spatial restrictions””. Stability points out that a “sole licence”

is a widely recognised type of intellectual property licence (see *Lufthansa Technik AG v Astronics Advanced Electronic Systems* [2025] EWHC 375 (Pat) at [636]-[637]) and that its case is supported by *JHP Ltd v BBC Worldwide Ltd* [2008] EWHC 757 (Ch), [2008] FSR 29 (“*JHP Ltd*”) in which a contract granting the “sole exclusive right” to publish a book was found to be an exclusive licence and not an assignment.

647. However, I agree with Getty Images that this is unsustainable. The employment contract does not need to assign any copyright to the entity employing Mr Rentz because as an employee any work created in the course of his employment vests in his employer in any event as a matter of law subject to any agreement to the contrary (see section 11 CDPA). That is the position under English law and Stability has not pleaded any principles of German law on which it seeks to rely to the contrary. Absent evidence of German law it is presumed to be the same as English law (see *Brownlie v FS Cairo (Nile Plaza) LLC* [2021] UKSC 45; [2022] AC 995 per Leggatt JSC at [115] and [143]-[149]). Neither of the authorities relied upon by Stability take matters any further. *JHP Ltd* was not in any event concerned with an employment relationship.

648. Further, on its true interpretation in context, I agree with Getty Images that clause 9(1) is neither the grant of a licence nor an assignment but is merely confirmation that Mr Rentz’s employer has the sole right to exploit all works arising out of his employment. It is not “an agreement to the contrary” for the purposes of section 11 CDPA. This is supported by the provisions of clause 11(1), which make clear that Mr Rentz does not own any images he creates, just as he is not licensing them to his employer:

“Upon termination of the employment relationship, Mr Rentz undertakes, for whatever reason, to hand over to the Company all business documents and documents in his possession as well as photographs (including transcripts and copies) and to confirm to the Company in writing that he has fully complied with this obligation. There is no right to retention of these documents”.

649. The only basis on which Mr Rentz could be required to hand over all of his work product upon termination is because it is owned by his employer. Mr Rentz himself appeared to me to confirm as much during his evidence (which I accept) when he confirmed that he had effectively “signed over” his rights to his employer. I reject Stability’s submissions in closing to the effect that clause 9(1) does no more than grant a perpetual licence.

650. I find that Getty Images has established its title to the copyright in SOCI Work A10 and A11.

(J) THE LICENSING ISSUE

651. A key issue in connection with the Secondary Infringement of Copyright Claim is whether the licensed Copyright Works relied on by Getty Images in these proceedings are subject to an exclusive licence within the meaning of section 92 CDPA in relation to which the First Claimant (or Second or Fifth Claimant) is the licensee. If those Copyright Works are subject to an exclusive licence, then Getty Images has concurrent rights with the copyright owners and is jointly entitled to a remedy for copyright infringement pursuant to sections 101 and 102 CDPA. If the licences are not exclusive, Getty Images lacks standing to bring a claim for copyright infringement in relation to

those works. Issues 5 and 6 of the Agreed List of Issues address the outstanding matters for determination by the court.

652. I have already addressed the procedural orders that were made with a view to facilitating the determination of the Licensing Issues. Fourteen Sample Licences have been identified with a view to determining the issues arising in respect of the licences identified in Annex 3A to the Particulars of Claim. Fortunately, the arguments on these licences may be distilled to a relatively small number of points, each based on identical or similar wording from licence to licence. A number of the Annex 3A licences are also relied upon to give Getty Images standing to bring claims in relation to SOCI Works A14, A15 and A16. By the time of closing submissions I understood it to be common ground that a finding on the Annex 3A licences can be read across to these SOCI Works.

Relevant Factual Background

653. I have already explained earlier in this judgment that Content is acquired by Getty Images in different ways, one of which involves Getty Images entering into licence agreements with third party photographers and copyright owners or “contributors”. Ms Cameron’s evidence, which I accept, is that this often involves photographers, as prospective contributors, submitting a sample of content via the Getty Images Websites and, following a review by the Getty Images’ Inspection Team, being offered either an iStock Artists Supply Agreement (“**iStock ASA**”) or a Unified Contributor Agreement (which I have already referred to as a Contributor Agreement). These are template agreements which have been updated and varied over time and which govern the terms on which the Content is licensed to Getty Images.

654. It is common ground that the terms of the Contributor Agreement and the iStock ASA are retrospective and cover all content previously submitted. Therefore, where a Contributor enters into a new version of either agreement, that version will supersede and cancel any previous agreement.

Relevant Legal Principles: English law

655. There is nothing between the parties as to the legal principles relevant in English law to the identification of an exclusive licence.

656. Section 92 of the CDPA defines an “exclusive licence” as follows:

“Exclusive licences.

(1) In this Part an “exclusive licence” means a licence in writing signed by or on behalf of the copyright owner authorising the licensee to the exclusion of all other persons, including the person granting the licence, to exercise a right which would otherwise be exercisable exclusively by the copyright owner.

(2) The licensee under an exclusive licence has the same rights against a successor in title who is bound by the licence as he has against the person granting the licence.”

657. Section 101 of the CDPA provides as follows:

“Rights and remedies of exclusive licensee.

- (1) An exclusive licensee has, except against the copyright owner, the same rights and remedies in respect of matters occurring after the grant of the licence as if the licence had been an assignment.
- (2) His rights and remedies are concurrent with those of the copyright owner; and references in the relevant provisions of this Part to the copyright owner shall be construed accordingly.
- (3) In an action brought by an exclusive licensee by virtue of this section a defendant may avail himself of any defence which would have been available to him if the action had been brought by the copyright owner.”

658. The essential requirements under s 92(1) are twofold: the licence is to be written and signed by the copyright owner and/or by someone on their behalf; and the licence is “to the exclusion of all other persons” including the copyright owner. The effect of sections 92(1) and 101 of the CDPA is that the exclusive licensee has all the rights of a copyright owner (i.e. as if there had been an assignment) and the copyright owner is left only with the right to sue other persons for infringement.

659. The relevant legal principles relating to exclusive licences were summarised in *Oxford Nanopore Technologies Ltd v Pacific Biosciences of California Inc* [2017] EWHC 3190 (Pat), [2018] Bus LR 353 by Mr David Stone sitting as a deputy High Court Judge at [44]. This was a patents case concerning the similarly worded definition of exclusive licence in the Patents Act 1977 sections 67 and 130(1). The parties are agreed that Mr Stone’s findings in [44] (set out below with reference to other authorities removed and with the wording tweaked to apply to a copyright case) are applicable to copyright law:

“(i) Whether or not a licence is an exclusive licence [under the relevant statute] is a matter for English law...;

(ii) A licence which purports to be an exclusive licence may not necessarily be so. Identifying an exclusive licence depends on a proper construction of the document or documents...An exclusive licence will be expressly so: circumstances in which an exclusive licence will be implied will be rare, if they exist at all;

(iii) It is for the party asserting that it is an exclusive licensee to demonstrate that it is...;

(iv) The assessment of whether or not a licence is exclusive is not a “once and for all assessment”...An exclusive licence may confer a power to convert the licence into a non-exclusive licence...;

- (v) The “essential element” of an exclusive licence is that it is a licence to the exclusion of all other persons, including [the copyright owner]...;
- (vi) It is possible to have a plurality of exclusive licences in respect of any one [copyright]...;
- (vii) But each exclusive licence may only be granted to one person – a licence will not be exclusive if granted to a number of entities, even if they are under the same control...;
- (viii) An exclusive licensee may grant sub-licences to “persons authorised by him”...;
- (ix) There is a distinction to be drawn between a licence and an equitable right to call for a licence...;
- (x) Where an equitable right to call for a licence is conditional... the otherwise exclusive licence will remain exclusive unless and until the contractual conditions are fulfilled that enable the grant of the licence...”.

660. Stability highlights point (v) which reflects a decision in *Dendron GmbH v University of California (No. 3)* [2004] EWHC 1163 (Ch) [2004] FSR 43 (“**Dendron**”), per Pumfrey J at [11]. *Dendron* was applied by Henry Carr J in *Illumina Inc v Premaitha Health Plc* [2017] EWHC 2930 (Pat). In the latter case the relevant licence purported to grant exclusive licence rights to a group of companies (as opposed to a single company). The learned Judge held (at [254]) that accordingly it was not an exclusive licence as “it is a licence granted by the proprietor to a number of persons, even though one of them is in *de facto* control of the others”.

661. Both parties emphasised that there may be a plurality of exclusive licences whose effect is to sub-divide the intellectual property right into separate exclusive domains, each relating to different rights, including different means of exploitation (see *Spring Form Inc v Toy Brokers Ltd* [2002] FSR 17 per Pumfrey J at [20] and *Copinger* at [4-247(2)] and [4-105]: “*The divisions which a copyright owner may desire to make will be into different modes of exploitation, different periods of time and different territories.*”

662. Where a licence agreement is governed by foreign law, the interpretation and construction of that agreement will be subject to its governing law. However, as is common ground, once any disputed issues of interpretation are resolved, the question of whether the licence is an “exclusive licence” is a question of English law applying sections 92 and 101(1) CDPA to the terms of the licence as the court has construed them. As Pumfrey J observed in *Dendron* at [9]:

“identifying an exclusive licence depends entirely upon a proper construction of the document or documents by which [the exclusive licensee] claims to be exclusive licensee”.

663. Two of the sample licences (both iStock ASAs, namely #33 and #36) are governed by the laws of Alberta, Canada. Getty Images’ pleaded case (with which Stability agrees)

is that the law of Alberta should be presumed to be the same as English law. The remaining sample licences (#2, #3, #8, #10, #11, #13, #17, #19 #30, #32, #37 and #38) are governed by the laws of the State of New York. I shall return to the principles to be applied to these latter agreements in a moment.

664. The principles of construction under English law are well known and I did not detect any disagreement between the parties as to the correct approach. Getty Images drew my attention to the court's summary of the relevant principles in *EE Ltd v Virgin Mobile Telecoms Ltd* [2023] EWHC 1989 (TCC) at [25]-[26]:

“25. i) The court construes the relevant words of a contract in their documentary, factual and commercial context, assessed in the light of (i) the natural and ordinary meaning of the provision being construed, (ii) any other relevant provisions of the contract being construed, (iii) the overall purpose of the provision being construed and the contract or order in which it is contained, (iv) the facts and circumstances known or assumed by the parties at the time that the document was executed, and (v) commercial common sense, but (vi) disregarding subjective evidence of any party’s intentions – see *Arnold v Britton* [2015] UKSC 36, [2016] 1 All ER 1, [2015] AC 1619 per Lord Neuberger PSC at paragraph 15 and the earlier cases he refers to in that paragraph;

ii) A court can only consider facts or circumstances known or reasonably available to both parties that existed at the time that the contract or order was made – see *Arnold v. Britton* (ibid.) per Lord Neuberger PSC at paragraph 20;

iii) In arriving at the true meaning and effect of a contract or order, the departure point in most cases will be the language used by the parties because (a) the parties have control over the language they use in a contract or consent order and (b) the parties must have been specifically focussing on the issue covered by the disputed clause or clauses when agreeing the wording of that provision – see *Arnold v. Britton* (ibid.) per Lord Neuberger PSC at paragraph 17;

iv) Where the parties have used unambiguous language, the court must apply it – see *Rainy Sky SA v Kookmin Bank* [2011] UKSC 50, [2012] 1 All ER (Comm) 1, [2012] 1 Lloyd’s Rep 34 per Lord Clarke JSC at paragraph 23;

v) Where the language used by the parties is unclear the court can properly depart from its natural meaning where the context suggests that an alternative meaning more accurately reflects what a reasonable person with the parties’ actual and presumed knowledge would conclude the parties had meant by the language they used but that does not justify the court searching for drafting infelicities in order to facilitate a departure from the natural meaning of the language used – see *Arnold v. Britton* (ibid.) per Lord Neuberger PSC at paragraph 18;

vi) If there are two possible constructions, the court is entitled to prefer the construction which is consistent with business common sense and to reject the other – see *Rainy Sky SA v. Kookmin Bank* (ibid.) per Lord Clarke JSC at paragraph 2 – but commercial common sense is relevant only to the extent of how matters would have been perceived by reasonable people in the position of the parties, as at the date that the contract was made – see *Arnold v. Britton* (ibid.) per Lord Neuberger PSC at paragraph 19;

vii) In striking a balance between the indications given by the language and those arising contextually, the court must consider the quality of drafting of the clause and the agreement in which it appears – see *Wood v Capita Insurance Services Ltd* [2017] UKSC 24, [2018] 1 All ER (Comm) 51, [2017] AC 1173 per Lord Hodge JSC at paragraph 11. Sophisticated, complex agreements drafted by skilled professionals are likely to be interpreted principally by textual analysis unless a provision lacks clarity or is apparently illogical or incoherent – see *Wood v Capita Insurance Services Ltd* (ibid.) per Lord Hodge JSC at paragraph 13; and

viii) A court should not reject the natural meaning of a provision as correct simply because it appears to be a very imprudent term for one of the parties to have agreed, even ignoring the benefit of wisdom of hindsight, because it is not the function of a court when interpreting an agreement to relieve a party from a bad bargain – see *Arnold v. Britton* (ibid.) per Lord Neuberger PSC at paragraph 20 and *Wood v. Capita Insurance Services Limited* (ibid.) per Lord Hodge JSC at paragraph 11”.

26. At [23] Sir Geoffrey Vos repeated the words of Lord Hodge in *Wood v Capita Insurance Services Ltd* [2017] UKSC 24 (at [12]), to the effect that the process of interpretation required is “a unitary exercise” (an observation also adopted by Lord Hamblen in *Sara & Hossein Holdings* at [29]), observing that:

“it starts with the words and the relevant context, and moves to an iterative process checking each suggested interpretation against the provisions of the contract and its commercial consequences. The court must consider the contract as a whole and give more or less weight to elements of the wider context in reaching its view as to its objective meaning”.”

Relevant Legal Principles: New York Law

665. I am extremely grateful to the parties for cooperating over the preparation of the Agreed Propositions of New York Law. I set these out below at paragraphs 666-675.

Sources of law

666. The principal source of New York state contract law are the decisions of the New York state courts including the New York Court of Appeals, the Appellate Division of the New York Supreme Court and the New York Supreme Court. Copyright is governed by federal law and in particular the Copyright Act of 1976 (“the Copyright Act”). However, a dispute over the interpretation of a contract to transfer the exclusive rights comprised in copyright is a state law matter (see *Gary Friedrich Enters., LLC v. Marvel Characters, Inc.*, 716 F.3d 302, 313).

The interpretation of written contracts

667. Contracts are interpreted in accordance with the parties’ intent (see *Greenfield v. Philles Records, Inc.*, 98 N.Y.2d 562, 569 (2002)). This is derived from the language employed in the contract. In seeking to ascertain the parties’ intent, the aim is a practical interpretation of the expressions of the parties to the end that there be a realisation of their reasonable expectations (*Harpercollins Publ., LLC v. Arnell*, 23 Misc 3d 1117(A) (Sup Ct 2009)). With written agreements, the best evidence of the parties’ intent is what is said in the written agreement (see *Slamow v. Del Col*, 79 N.Y.2d 1016, 1018). Thus, when parties set down their agreement in a clear, complete document, their writing should as a rule be enforced according to its terms (*W.W.W. Assocs. v. Giancontieri*, 77 N.Y.2d 157, 162 (1990)).

668. Written agreements are interpreted as a whole and each term is interpreted in the light of the rest of the agreement, with all writings that are part of the same transaction interpreted together. Particular words should be considered, not as if isolated from the context, but in the light of the obligation as a whole and the intention of the parties as manifested thereby. Form should not prevail over substance and a sensible meaning of words should be sought (see *S & S Media, Inc. v. Vango Media, Inc.*, 84 A.D.2d 356, 359 and Restatement (Second) of Contracts §202(1) and (2)). Unless a different intention is manifested, where language has a generally prevailing meaning, it is interpreted in accordance with that meaning and where technical terms or words of art are used, these are given their technical meaning (see Restatement (Second) of Contracts §202(3)(a) and (b)).

669. In construing a contract, the Court “should arrive at a construction that will give fair meaning to all of the language employed by the parties to reach a practical interpretation of the expressions of the parties so that their reasonable expectations will be realized” (see *Petracca v. Petracca*, 302 AD2d 576, 576-77 (2d Dep’t 2003)). The contract must be read as a whole to determine the parties’ purpose and intent (*Krape v. PDK Labs, Inc.*, 34 A.D.3d 751, 753).

670. Where several constructions of a contract are possible, the court can look to the surrounding facts and circumstances to determine the intent of the parties (see *25 Bay Terrace Associates, L.P. v. Public Service Mutual Insurance Company*, 194 A.D.3d 668, 670). The background for this purpose includes “the entire situation, as it appeared to the parties” at the time of entering the contract (Restatement (Second) of Contract §202 at (b)). Where a contract is ambiguous, extrinsic evidence may be considered to determine the parties’ intent (*Yarom v. Poliform S.P.A.*, 153 AD3d 760, 761 (2d Dep’t 2017)).

671. A contract should not be interpreted in such a way as to leave one of its provisions substantially without force or effect but, instead, to “give fair meaning to all of the language employed by the parties to reach a practical interpretation of the expressions of the parties so that their reasonable expectations will be realized” (*Petracca v. Petracca*, 302 AD2d 576, 576-77 (2d Dep’t 2003)). The court must avoid an interpretation that would leave contractual clauses meaningless and/or would render contractual language mere surplusage (*Two Guys from Harrison–N.Y., Inc. v. S.F.R. Realty Associates*, 63 N.Y.2d 396, 403 (1984); see also *Westview Assoc. v. Guaranty Natl. Ins. Co.*, 95 N.Y.2d 334, 339 (2000)).

672. A relevant aspect of the context to an exclusive licence agreement over copyright is the federal Copyright Act which governs issues relating to the ownership and licensing of copyright and to the issue of standing to sue and which has been interpreted in decisions of the federal District Courts, the US Circuit Courts of Appeal and the US Supreme Court.

673. Normal principles of contractual interpretation apply to agreements to transfer copyright (*Orange Cnty. Choppers, Inc. v. Olaes Enterprises, Inc.*, 497 F. Supp. 2d 541, 551 (S.D.N.Y. 2007)).

674. Under §101 of the Copyright Act, a “transfer of copyright ownership” is defined to include “an exclusive licence...of a copyright or of any of the exclusive rights comprised in a copyright whether or not it is limited in time or place of effect, but not including a nonexclusive licence”. Under §501(b) of the Copyright Act, “the legal or beneficial owner of an exclusive right under a copyright is entitled...to institute an action for any infringement of that particular right”. That includes an exclusive licensee within the meaning of §101 of the Copyright Act (*Minden Pictures, Inc. v. John Wiley & Sons, Inc.*, 795 F.3d 997, 1004 (9th Cir. 2015) (“**Minden Pictures**”)).

675. A licence agreement is exclusive where it permits a licensee to use the protected material for a specific use and further promises that the same permission will not be given to others (see *I.A.E., Inc. v. Shaver*, 74 F.3d 768, 775). The exclusive rights subject to an exclusive licence may be shared between one or more parties, including between the licensee and the owner and between more than one licensee, without the licence falling outside of §101 of the Copyright Act (see *Minden Pictures* at 1004).

676. The latter proposition of New York law is of course very different from the position under the CDPA referred to above, where an exclusive licence comprising the same rights cannot be granted to multiple entities, or, put another way, a licence will not be exclusive if it is not exclusive to one party. In brief summary, it is Stability’s contention (to which I shall return) that this point is of importance when it comes to construing the licence agreements governed by New York law which it says appear to have been drafted with this more permissive regime in mind. Stability submits that the use of the word “exclusive” in those licences is use of a legal term of art, or a technical term, and that it is likely to be a reference to the New York law concept of an exclusive licence (which may include multiple parties) rather than an allusion to any English law concept of an exclusive licence (which can involve only one party).

Analysis of the outstanding issues of construction

677. In broad terms, Stability challenges the exclusivity of the relevant licence agreements on three main grounds. These have been referred to by the parties as (i) “the Multiple Entity Point”; (ii) “the Carve Out Point”; and (iii) “the Signature Point”. An additional point identified at trial by the parties as “the Date Point” had fallen away by the time of closing submissions. I shall address these points in turn.

(i) The Multiple Entity Point

678. This point concerns only the Contributor Agreements (thereby bringing into focus the principles of New York law to which I have referred). Although it arises slightly differently in relation to variations of these agreements owing to variances in their terms, Stability’s fundamental point is that the relevant licences have been granted to more than one entity (as permitted by New York law) such that, as a consequence, they cannot be exclusive licences for the purposes of English law. Getty Images would have standing to sue in the United States, but they do not have standing to do so here.

Sample Agreement #2

679. I start by focusing (as did Getty Images in oral submissions) on sample licence #2. This agreement dates from February 2007. Its signature page identified the Contracting Parties, the Contributor and “Getty Images (US) Inc., a New York corporation with offices located at 75 Varick Street, New York, NY 10013, U.S.A. (“**Getty Images**” or “**we**””).

680. A single recital then records that:

“Getty Images (US), Inc. (“**Getty Images**” or “**we**”) invites the Contributor identified on the signature page to the agreement...to submit photographic images, illustration, audio, footage, fonts and/or animation that may be used in Getty Images’ products and services and for licensing and distribution by Getty Images to third parties. The following terms and conditions will apply to all material submitted and accepted for licensing by Getty Images under this Agreement”.

681. The licence provision is at clause 2.1 and is in the following terms:

“2.1 License Grant to Getty Images. Subject to Section 2.3, you grant Getty Images a worldwide, exclusive right to distribute, market, sublicense, rent, use, copy, reproduce, publish, transmit, broadcast, display, communicate and make Accepted Images available to the public...Getty Images may sublicense or authorize any of its Distributors, Clients and their customers to exercise the rights described in this paragraph...”.

682. Section 1 to the agreement is a Definitions section. “**Getty Images**” is defined as meaning:

“Getty Images (US), Inc and each of the entities, controlling, controlled by or under common control with Getty Images”.

In the same section, “Distributor” is defined as meaning “any company, firm other organization, entity or person authorized by Getty Images to license the Accepted Images to Clients”. “Accepted Images” are images “accepted for distribution under the Agreement”.

683. Against that background, Stability submits that, on its true construction under New York law, agreement #2 plainly grants licence rights to multiple parties, namely multiple companies in the Getty Images Group. Getty Images submit, on the other hand, that there are two definitions of Getty Images in this agreement – the definition appearing in the recital and the definition appearing in section 1. They say that the definition in the recital is clear and unambiguous whereas the definition in section 1 is unclear, confused and circular because “the characteristic that links each of the entities identified in the list (i.e. those controlling, controlled by, or under common control) is their relationship with “Getty Images”. This, it is submitted, simply raises the question, what is “Getty Images”? It would make no commercial sense to favour the definition in section 1 over the very first sentence of the agreement. Further, says Getty Images, if the definition in section 1 were to be adopted, then the definition in the first sentence of the agreement “is entirely redundant” – an outcome which would be inconsistent with the approach taken to construction under New York Law.
684. During oral submissions, Ms Bowhill drew my attention to three specific points which she said supported the interpretation for which Getty Images contends:
 - i) First, she points to the fact that clause 2.1 grants a “worldwide exclusive right”, submitting that this must have been intended to cover the United Kingdom. As I understand it, Getty Images’ point is that there would have been a reasonable expectation to this effect.
 - ii) Second she points to the provisions of clause 2.1 in relation to the ability of Getty Images to sublicense or authorise its Distributors to exercise the rights granted by that clause; she contends that, on its face, the definition of “Distributor” includes other Getty Images entities such that the Accepted Content is sub-licensed to those entities as Distributors - it is not licensed directly to the Getty Images Group as a whole. Ms Bowhill submits that if the Getty Images entities are already Distributors, it would be superfluous for them to be included in the definition of “Getty Images”.
 - iii) Third, Ms Bowhill relied upon the court’s entitlement to look at the surrounding facts and circumstances to determine the intent of the parties and she pointed to evidence of the First Claimant sub-licensing its rights to other entities in the Getty Images Group, such as the Second Claimant.
685. Having regard to the principles of New York law to which I have been referred, I prefer the construction placed on #2 by Stability for the following main reasons:
 - i) I start from the proposition that the parties plainly intended to define the meaning of the words “Getty Images” and that they did so in a “Definitions” section at section 1. While it is true that there is potential ambiguity in the

definition caused by its apparent circularity, the use of the words “and each of the other entities, controlling, controlled by or under common control with” must have been intended by the parties to mean something. On Getty Images’ construction, these words are meaningless and amount to no more than contractual surplusage. New York law requires the court to avoid any construction that would render a contractual provision meaningless.

- ii) Given that the parties have expressly sought to define the term “Getty Images” it is difficult to see why they would have regarded the content of a recital as a “definition”, as Getty Images contend. In any event, in my judgment the recital is plainly not intended as a definition. Instead it is an invitation by one of the parties to the agreement to the other (as identified on the signature page) to submit photographic images together with a statement that the terms and conditions set out thereunder will apply to such arrangement. When the recital is seen in these terms, there is nothing inconsistent between the content of the recital and the definition section.
- iii) As for the potential ambiguity in the definition section relating to “Getty Images”, it appears to me to be important to have regard to the need to give a fair meaning to all of the language employed by the parties with a view to reaching a “practical interpretation”. As Stability suggests, a practical interpretation (which gives effect to the remaining wording in the provision) is that the second use of Getty Images in the definition in section 1 is intended as a reference to Getty Images (US), Inc or the general Getty Images Group. On either reading, the definition of Getty Images involves multiple companies in that it refers to all companies in common ownership within the Getty Images Group. On a practical interpretation of this definition, there is no real ambiguity.
- iv) I do not regard the reference to “worldwide exclusive rights” to be persuasive. I must have regard to the intention of the parties in circumstances where they were contracting at all times against the background of New York law, which permits exclusivity even where there are multiple licensees; I disagree with Ms Bowhill that it is practical or realistic to construe the agreement on the assumption that the parties were intending to ensure exclusivity in accordance with the laws of the UK. As Mr Edwards submits, the licence is exclusive in the United Kingdom in the primary sense of that term because the photographer/Contributor is not permitted to license the UK copyright to any company outside the Getty Images Group under the terms of this agreement. The Contributor Agreement is an exclusive agreement in the ordinary sense of that term, even on Stability’s construction.
- v) I also do not consider Ms Bowhill’s arguments as to the definition of “Distributors” to assist, not least because Ms Bowhill’s submissions on the point assume that “Getty Images” means only Getty Images (US), Inc. On that assumption, it is true that the definition of “Distributor” could cover other Getty Images entities, although one wonders why the parties referred to them in the terms they did in the definition of “Distributor” when they had expressly given consideration only a few lines earlier to “entities, controlling, controlled by or under common control with”. When read together with the definition of “Getty Images” and bearing in mind the very different language used, it appears to me to be clear that the two definitions are intended to deal with different things.

There is nothing inconsistent or superfluous in a reading of the definitions which treats them as intending to address, on the one hand Getty Images entities (within the definition of Getty Images) and on the other hand any company, firm, other organization, entity or person other than Getty Images entities (within the definition of “Distributor”).

- vi) Ms Bowhill pointed to clause 4.7 of agreement #2 (which provides that the Contributor authorises “Getty Images and Distributors at their expense to make, control, settle and defend Claims”), submitting that the reference here to Getty Images can only be to Getty Images (US), Inc alone “because otherwise every affiliate entity would be being given the right to bring and control proceedings” anywhere in the world. She submitted that this would create “total chaos” because it would mean that no specific entity within the Getty Images Group would have ultimate control, including over proceedings in the UK.
- vii) However, the problem with this argument, as Mr Edwards pointed out, is that clause 4.7 plainly gives “Distributors” the right to “make, control, settle and defend Claims”, thereby already relinquishing the prospect of any one corporate entity having ultimate control. The right to control claims is plainly not envisaged to depend on the First Claimant’s standing in its own right to bring claims (or, indeed, any other Getty Images entity). While Getty Images contend that the definition of Distributors must include Getty Images entities (a submission which I have rejected), it nevertheless accepts that Distributors must also include third party entities. Such third party Distributors would have no standing under the Contributor Agreement as exclusive licensees or otherwise to bring any claim in the UK and yet this clause gives them the right to control claims and requires the Contributor to cooperate with any such claims. In a similar way, the First Claimant and other Getty entities will have the right to control claims. I agree with Stability that the reasonable expectations of the First Claimant are fulfilled where it has the right to bring proceedings in the Contributor’s name in any event.
- viii) Finally, given the approach I have already adopted, I reject the suggestion that there is ambiguity which requires the court to look at the surrounding facts and circumstances to determine the intent of the parties. However, even if I am wrong about that, the evidence to which my attention was drawn by Getty Images does not assist their case. Ms Bowhill relied upon the unchallenged evidence of Mr Prichard to the effect that inter-company arrangements provide for sub-licensing to other entities in the Group. She did not suggest that this was evidence of which Contributors would have been aware or explain how the court could find that it was relevant to “the entire situation as it appeared to **the parties**” (emphasis added), particularly in circumstances where much of it was said to be confidential in the context of the proceedings. There is no evidence on which I could properly make such a finding. Mr Edwards also pointed out in submissions that Mr Prichard’s evidence all post-dates Sample Licence #2 because he was not employed by the Getty Images Group until 2012 and the documents to which he refers in his evidence all date from later than 2015.

686. For all the reasons set out above, I find that Sample Agreement #2 is not an exclusive licence under section 92 CDPA.

Sample Agreements #3, 8, 10, 11, 13, 30 and 32

687. These agreements cover the period between July 2007 and October 2011. The key material difference between this group of agreements and Sample Agreement #2 is that Getty Images is defined in these agreements in the following terms (key differences appear in bold):

- i) Sample Licence #3:

“Getty Images (US), Inc and, **where the context infers**, each of the entities, controlling, controlled by or under common control with **Getty Images (US), Inc.**”.
- ii) Sample Licences #8, 10, 11, 13 and 30:

“Getty Images (US), Inc and, **where the context implies**, each of the entities, controlling, controlled by or under common control with **Getty Images (US), Inc.**”.
- iii) Sample Licence #32:

“[Getty Images (US), Inc.]/[**Getty Images International**] and, **where the context implies**, each of the entities, controlling, controlled or under common control with [**Getty Images (US), Inc.**]/[**Getty Images International**]”.

688. Getty Images accept that these definitions all make it clear that, in certain contexts throughout the relevant Contributor Agreements, the term Getty Images means Getty Images (US), Inc and its affiliated companies. However, they contend that the obvious (and stated) intention is that there will be contexts in which that term means only Getty Images (US), Inc. Thus the term may have a different meaning depending on which clause is in issue and the context of that clause.

689. I observe that Sample Licence #3 is the only licence in this group that also has wording in its recital which matches the wording in the recital to Sample Licence #2. However, given that the definition of “Getty Images” in Sample Licence #3 could not possibly be said to be ambiguous, confusing or circular, I cannot see that this affects the analysis to be conducted in relation to Sample Licence #3. I did not understand Getty Images to suggest otherwise.

690. The parties are agreed that the wording in the grant of rights clause in these Sample Licences is broadly consistent across all relevant agreements and I did not understand them to submit that the slightly different wording of other terms and conditions in each agreement affected the overall analysis. Broadly, then, the question for the court is how the words “Getty Images” used in the grant of rights clause are to be construed.

691. I take the grant of rights clause in the “Standard Terms and Conditions” of Sample Licence #8, as Getty Images did in their submissions, by way of example¹². As I understood the parties submissions, the points raised in relation to Sample Licence #8

¹² In this group of Agreements the grant of rights clauses are all at 1.1, save for Sample Licence #3, where it is to be found in clause 2.1).

apply equally to agreements #3, 10, 11, 13, 30 and 32. The grant of rights clause in #8 reads as follows:

“1.1 License Grant to Getty Images: You grant Getty Images a worldwide, exclusive right to market and sublicense Reproduction Rights in Accepted Content. Except for Still Images that you submit for licensing through Rights-Managed Editorial or the Reportage Collection (“Reportage”), you also grant Getty Images the additional right to modify, adapt or create Derivative Works of all other Content, in any medium and for any purpose. Getty Images may sublicense or authorize any Distributors, Clients and their customers to exercise the rights described in this Section 1. Getty Images and Distributors will determine the terms and conditions of all licenses of Accepted Content granted by them, but will not use or license Accepted Content for uses that are defamatory, pornographic or otherwise illegal and will use commercially reasonable efforts to stop any such use brought to our attention. Getty Images and Distributors may determine how Accepted Content may be marketed and may stop marketing or licensing it at any time. If Getty Images notifies you that it has permanently stopped marketing and licensing Accepted Content, the Agreement will be deemed to be terminated only with regards to that Accepted Content”.

692. As Getty Images point out, Sample Licence #8 makes clear in its “Commercial Terms” under the heading “Exclusivity” (“**the Exclusivity Provision**”) that:

“[a]ll Content submitted to Getty Images is on a Content exclusive basis. Once Content has been submitted to Getty Images, such Content and any Similars may not be licensed to any third party unless Getty Images has notified you that it has been rejected. In addition, you must submit exclusively to Getty Images any Content or Similars that you have created (a) on assignment for or as representative of Getty Images; (b) acting on information, direction or access provided through Getty Images; or (c) where Getty Images is funding any of the costs incurred in connection with the creating that content”.

693. A “Right to Control Claims” clause at 1.4 of the standard terms and conditions provides that:

“You authorize Getty Images and Distributors at their expense the exclusive right to make, control, settle and defend Claims related to Accepted Content... You agree to provide reasonable cooperation to Getty Images and Distributors...”

694. In the definitions section at section 5 of Sample Licence #8, “Distributor” is defined as “any company, person, or other entity **authorized** by Getty Images, directly or indirectly, to license Accepted Content to Clients” (**emphasis added**). “Client” means “any customer who licenses Accepted Content from Getty Images or a Distributor”.

695. Once again, Stability submits that, having regard to the context, the reference to Getty Images in the grant of rights clause is a reference to multiple parties with the effect that it is not an exclusive licence under section 92 CDPA (“**the Wider Construction**”).

696. Getty Images contend, on the other hand, that the grant of rights clause at 1.1 is intended to grant rights only to Getty Images (US), Inc (“**the Narrow Construction**”). However, they accept that “several constructions are possible” and accordingly submit that the court is permitted to look at the surrounding facts and circumstances to determine the intent of the parties. Specifically, Getty Images submit that relevant surrounding facts and circumstances include that:

- i) the entity within the Getty Images Group that licenses Copyright Works to end consumers in the US is the First Claimant, Getty Images (US), Inc; and
- ii) the entity within the Getty Images Group that sub-licenses Copyright Works to the Second Claimant (who sub-licenses Copyright Works to related entities outside the US) is the First Claimant, Getty Images (US), Inc.

697. Thus Getty Images say that a practical interpretation of clause 1.1 points to the term “Getty Images” meaning only Getty Images (US), Inc. They submit that it makes commercial sense that the First Claimant is the only entity to benefit from the grant of the exclusive licence because it is the entity that (i) enters into the agreement with the Contributor; (ii) can enforce the agreement if it needs to; and (iii) sub-licenses the Copyright Work thereafter. If the exclusive licence being granted extended to all of the First Claimant’s affiliated companies, it would be pointless for the First Claimant to then grant sub-licences to other entities in the Group, such as the Second Claimant. Thus, say Getty Images, the fact that the First Claimant grants such sub-licences supports the proposition that the parties to this Contributor Agreement intended to grant an exclusive licence to a single entity only, i.e. the First Claimant.

698. Getty Images also rely upon the submission (as recorded in relation to Sample Licence #2 above) that the definition of “Distributors” covers other Getty Images entities and so renders it unnecessary for an exclusive licence to be granted more widely. They contend that the Narrow Construction is consistent with *Petracca v Petracca* in that it gives a fair meaning to all of the language employed by the parties to reach a practical interpretation of the expressions of the parties so that their reasonable expectations will be realized. Specifically, it is (say Getty Images) the only construction which (i) gives exclusivity to the licence in the UK; and (ii) allows the First Claimant to control claims in the UK “which cannot occur if Getty Images in clause 1 is construed to cover the First Claimant and its affiliated companies”.

699. Getty Images point to other clauses in Sample Agreement #8 which they say either support (in context):

- i) the Narrow Construction: e.g. the “Representation and Warranties” clause at 2.1; a proposition with which Stability agrees; or
- ii) the Wider Construction: e.g. some (but not all) of the references to “Getty Images” in the Exclusivity Provision; clause 2.2 which provides for “Indemnification”; and clause 4.8 which provides for “Confidentiality”).

700. Once again, I prefer the Wider Construction of the grant of rights clause in Sample Licence #8 for the following reasons:

- i) I cannot see how I could properly find that the facts and circumstances on which Getty Images rely are part of “the entire situation, as it appeared to the parties” at the time of entering the contract. I note in this regard that the *Restat 2d of Contracts* §202 makes clear that in appropriate cases “the entire situation, as it appeared to the parties”, may include “facts known to one party **of which the other had reason to know**” (**emphasis added**). In support of the factual background which I have set out above and on which they rely, Getty Images point to various paragraphs in the unchallenged evidence of Mr Prichard (albeit again his direct knowledge appears to post-date these agreements) and Ms O’Neill (who was not employed by the Getty Images Group until late 2015) which address the relationships between various Getty Images entities including the practice of sub-licensing. However, aside from the question whether these witnesses have relevant direct evidence to give, all or much of this evidence appears to be subject to confidentiality restrictions in these proceedings and (even if this information was not said to be confidential) it is, in any event, wholly unclear how it is said that Getty Images’ internal commercial arrangements which have clearly changed over time would be known to a Contributor entering into a Contributor Agreement at any particular juncture. There is no evidence whatever addressed to the issue of what information a Contributor would have had reason to know when signing up to a Contributor Agreement.
- ii) Absent any other relevant surrounding facts (and Getty Images did not point me to anything else), I am left to construe the wording of the Sample Licence itself, having regard to the principles of New York Law.
- iii) Looking first at the Exclusivity Provision (which is concerned with the party or parties which benefit from the exclusive licence rights granted under the Contributor Agreement), I reject the submission made by Getty Images that the reference to “Getty Images” means something different in different parts of this provision. I agree with Stability that such an interpretation would be both impractical and confusing, leaving the reader to assign different meanings to different usages within the same clause. Absent a very clear and express contrast between one usage and another (made clear with appropriate clarificatory language), Getty Images’ approach would not, in my judgment, accord with the need to give fair meaning to the language to reach a “practical interpretation”. A practical interpretation of the Exclusivity Provision would assign the same meaning to all usages within the same clause. If, as Getty Images submit, one of the references to Getty Images in this clause is plainly a reference to the First Claimant and its affiliates, then, to my mind, that context makes it more likely than not that the parties intended all other references to “Getty Images” in the provision to mean the same thing.
- iv) Standing back, and having regard (a) to the fact that the definition of Getty Images appears to envisage the exploitation of licensed content by the Getty Images Group as a whole; and (b) the potential under New York law for an exclusive licence to be shared between multiple parties, I consider that an interpretation which reads every reference to “Getty Images” in the Exclusivity

Provision as a reference to the First Claimant and its affiliates is most likely to realise the “reasonable expectations” of the parties.

- v) Given that, on balance, it appears to me that the Wider Construction applies to the entirety of the Exclusivity Provision, I accept Stability’s submission that it must equally apply to the grant of rights clause at 1.1. I reject Getty Images’ submission that this is inconsistent with the grant of a “worldwide, exclusive right” or with the need for the First Claimant to control claims, for reasons I have already given. In the circumstances there seems to me to be nothing in the context of clause 1.1 to support the Narrow Construction, which I do not consider to be inconsistent with the guidance in *Petracca v Petracca*.
- vi) I am not persuaded that the meaning of the defined term “Distributors” undermines Stability’s submissions. On a natural reading of that defined term in Sample Licence #8 (and #3, 10 and 30 which all defined Distributors in the same terms), I consider it plainly to be intended to mean “any company, person or other entity” outside the Getty Images Group. Again, the context of this definition, so close to the definition of “Getty Images” which uses entirely different language to refer to entities within the Getty Images Group, appears to me to make plain the clear distinction between the entities covered by each definition. In defining “Distributor” the parties have expressly chosen to refer to the concept of entities “authorized by” Getty Images, whereas a few lines later they have talked about entities within the Getty Group using the language of control. The “Personal Data Transfer” clause 4.9 of Agreements #8 and 10 (which expressly distinguishes between “related companies” and “Distributors”) appears to me to support this construction.
- vii) A similar conclusion is inevitable, in my judgment, in light of the slightly differently worded definition of “Distributor” in Sample Licences #11, 13 and 32 (embedded within clause 1.1) which provides that:

“Getty Images may sublince or authorize **any third party** distributors (“**Distributors**”), any customer who licences Accepted Content from Getty Images or a Distributor (“**Clients**”) and their customers to exercise the rights described in this Section 1” (**emphasis added**”).

It is clear from this definition that the Distributor must be a third party; in other words an organisation that is not affiliated with Getty Images. Again this interpretation is supported by the “Personal Data Transfer” provisions in these agreements (at clause 4.9 in #11 and #13 and at clause 4.8 in #32) which distinguish between “related companies” and “Distributors”.

701. For all the reasons set out above, I find that Sample Licences #3, 10, 11, 13, 30 and 32 are not exclusive licences for the purposes of section 92 CDPA.
702. I should, however, add a postscript. It appeared to me, when considering the construction of this group of agreements, that an argument that could have been advanced by Getty Images was to the effect that the wording in the various definition sections made clear that the *default position* when considering the meaning of the words “Getty Images” was always that those words mean Getty Images (US), Inc. unless the

context leads to the inevitable inference that those words mean the Group companies (i.e. “and, where the context implies/infers”). It is not clear to me that this point would necessarily have made any difference (particularly in light of my decision above in relation to the interpretation of the Exclusivity Provision), but, in any event, as it was not argued by either side, I do not consider it appropriate to address it further.

Sample Licences #17 and #19

703. Finally, I must turn to consider Sample Licences #17 and 19 which cover the period June 2014 to December 2018. Sample Licence #19 is the latest agreement in issue in these proceedings.

704. All references to “Getty Images” in Sample Licences #17 and #19 are said (at the outset of those agreements) to mean:

“Getty Images (US), Inc and, where the context implies, each of the entities, controlling, controlled or under common control with Getty Images (US), Inc (each an “Affiliate”). The rights granted to Getty Images under this Agreement may be sublicensed to one or more Affiliates in Getty Images’ discretion”

705. Does this alter the analysis to which I have referred above? In particular, perhaps, what is the impact of this provision when read in conjunction with an Exclusivity Provision (in much the same terms as the Exclusivity Provision referred to above) in which Getty Images accepts that at least one of the references to Getty Images must be to the wider group of companies, but says that other references to Getty Images in the same provision are to be interpreted narrowly?

706. I have already held that Getty Images’ approach to the Exclusivity Provision is confusing and impractical for the purposes of the argument on interpretation in relation to the earlier Sample Licences, but how should the court approach the Exclusivity Provision in these later Sample Licences where the definition of Getty Images is in very different terms?

707. For reasons I have explained, I have no extraneous factual background or surrounding circumstances on which I can rely to provide me with evidence of the parties’ shared intent and reasonable expectations. However, Getty Images submits that there can be no real doubt that if the grant of rights clause in 1.1 of the standard terms and conditions of these two agreements was intended to grant rights to Getty Images Group there would have been no need to make provision for those rights to be “sublicensed to one or more Affiliates in Getty Images’ discretion”. Those words would be meaningless.

708. In response, Stability submits that the additional wording in the definition gives Getty Images “the option” to sub-license to Affiliates and that “if it does not exercise this option, Getty needs the licence to cover the whole Group because all its companies are involved in exploiting this copyright”. It relies on the same arguments in relation to the Exclusivity Provision and it contends that Getty Images’ consistent practice of taking a licence which grants rights to the whole Group stretches right back to its early Contributor Agreements and that it is to be inferred that Getty Images do not intend to agree terms in these later agreements which change that practice.

709. On balance, however, and having regard to the need to avoid an interpretation that would render contractual language “mere surplusage”, it seems to me that the definition of Getty Images in the grant of rights clause in these agreements (and in parts of the Exclusivity Provision dealing with content being submitted to Getty Images), must be given the Narrow Construction. This is notwithstanding that this involves interpreting different references to Getty Images in the Exclusivity Provision in different ways, without any appropriate clarificatory language. While this is to my mind undesirable and while clear words would ordinarily be required to render this necessary, I am left with the conflicting indications provided by the Exclusivity Provision and the definition of Getty Images and I consider that the words now included in the definition of “Getty Images” must carry the day if the principles of New York law to which I have been referred are to be applied. Stability’s explanation that the additional words provide Getty Images with “an option” do not, to my mind, address the fundamental issue that there is no need for that option if the licence is to the Group as a whole. Stability provided me with no other theory as to how the additional wording in the definition could be given a meaning if it was right in its interpretation.

710. The fact that the definition of “Distributor” in these agreements makes reference to “third party distributor” – thereby to my mind clearly making the distinction between a Distributor on the one hand and an Affiliate on the other – does not appear to me to alter this conclusion. Equally, in the face of the additional words used by the parties when defining Getty Images, I do not consider the fact that New York law permits exclusive licences to be made with more than one licensee to be persuasive.

711. I reject Stability’s submission that the evidence of the earlier Sample Licences indicates Getty Images’ consistent practice and that these later agreements should be interpreted in that same light. Where wording has been changed, it is difficult to see how any presumption of “consistent practice” could properly be applied. In any event, I was shown no principle of New York law to this effect. I note that each agreement is subject to an entire agreement clause making clear that any previous agreements related to the distribution of Accepted Content have been superseded and cancelled and that the agreement “constitutes the entire agreement among the Parties relating to its subject matter...”.

712. Finally, Stability submitted that there is no evidence of the First Claimant granting sub-licences to Getty Images affiliates despite a large number of Getty companies being involved in the marketing and sale of images and despite the fact that Getty Images have disclosed large numbers of other agreements between different Getty group companies as part of their case on title. Stability suggests that the inference to be drawn is that “Getty rely principally on the head licence rights granted by these agreements to exploit these images as these rights extend [beyond the First Claimant] to the rest of the Getty group”. I reject this submission. I have seen nothing in the agreed principles of New York law which would permit me to take this course and nor did Stability draw my attention to any principles on which it relied in making this submission.

713. I do not consider any of the other arguments canvassed above in respect of the other Sample Licences to alter this outcome. The references to Getty Images in the various clauses of these agreements must be construed having regard to the context of the definition to which I have referred.

714. Accordingly I find that on their true construction Sample Licences #17 and #19 are exclusive licences pursuant to section 92 CDPA.

(ii) The Carve Out Point

715. As developed in closing, I understood this point to concern only the iStock ASAs, whose construction, as I have said, is to be determined having regard to the canons of construction in English law.

716. Each of the iStock ASAs in Sample Licences #34-38 contain a provision in the following terms at clause 9(b):

“You represent and warrant that you shall not: (i) license your own Exclusive Content (**except occasionally and then only for legitimate creative purposes**); or (ii) predominately license the content of only a few contributors. You agree that you will not collude with another iStockphoto member to have that member do either of (i) or (ii) above for your benefit. You acknowledge that genuine subscription customers typically license files from many contributors and you agree that your subscription licensing behavior will conform to this typical conduct. In addition to any other available remedies, if you breach this paragraph iStockphoto may immediately terminate this Agreement and/or, if applicable, cancel your subscription package without any refund to you. You further agree to forfeit any royalties earned by you in connection with your misconduct.” (**emphasis added**).

717. Stability submits that what it describes in its Defence as a “vaguely defined carve-out” appears to permit each copyright owner to license his works and exploit the Content in parallel with the First Claimant (and/or as the case may be, its predecessor licensees). It says that this is therefore not a license in the First Claimant’s favour “to the exclusion of all other persons, including the person granting the licence”. The copyright owner is permitted to license his content as long as this is only done “occasionally” and for “legitimate creative purposes”. Stability says these are qualitative limitations that do not limit the copyright owner’s rights in this respect “to any specific defined area”.

718. Getty Images contend, on the other hand, that this is not a carve-out of rights at all. They point to the licensing provisions in these agreements which they say expressly reserve exclusive rights to Getty Images and they submit that the parties cannot have intended that these be overridden by clause 9(b), which (they say) is no more than a representation and warranty from the licensor which simply contains “a caveat” in the form of the emboldened words identified above. Thus, say Getty Images, clause 9(b) does not permit the licensor to do something, it simply acknowledges that, if the licensor does a specific purpose-limited act, it will not be in breach of its representation or warranty. In closing oral submissions, Ms Bowhill also submitted that the provisions of clause 9(b) are plainly a reference back to clause 2(b), which is where the rights which are retained by the licensor are set out. She emphasised that on no sensible construction could the words of this caveat include a carve out in respect of the licensing of images for training an AI model.

719. Taking Sample Licence #36 by way of example (as Getty Images did in oral submissions), clause 1(c) provides that:

“The supplier wishes to appoint iStockphoto as its exclusive agent to license, sublicense and distribute Exclusive Content (as defined below) produced by the Supplier on the terms and conditions set forth in this Agreement”.

720. Clause 2(a) defines the Exclusive Content as one or more of “(i) Photo Content, (ii) Illustration Content; (iii) Flash Content or (iv) Motion Content”. Clause 2(b), provides that:

“Notwithstanding the definition of Exclusive Content and the exclusive license granted in this Agreement, nothing shall restrict the Supplier from (i) establishing or maintaining a personal portfolio web-site on which Exclusive Content is posted for the purposes of art display **but not the sale or licensing or giving away of rights to the digital Content**; or (ii) using Exclusive Content in connection with the sale by Supplier of prints, t-shirts and other merchandise **where the sale or licensing or giving away of rights to the digital images or other Content beyond such merchandising use is not involved.**” (**emphasis added**)

721. Pausing there, it is clear that clause 2(b) expressly prohibits a licensor from selling, licensing or giving away rights to the digital Content save for the purposes of selling merchandise. As Getty Images submits, all rights beyond those reserved to the Supplier in clause 2(b) are expressly granted to Getty Images and are exclusive. This is clear from the wording of clauses 1(c) and 2(b). I note in this context that this agreement is identified as an “Artist’s Supply Agreement (Exclusive)” and that the recital records that “[t]his Agreement governs the terms by which photographers, videographers or other artists provide stock photographic, video and other media content to members of the iStockphoto.com community, on an exclusive basis through web site located at www.istockphoto.com (the “iStock Site”), and to other prospective purchasers through other distribution venues as provided for in this Agreement. [For the non-exclusive Artist’s Supply Agreement, go to this link]” (original emphasis).

722. It is very difficult to see how a clause dealing with representations and warranties can have been intended to undermine this purpose. Furthermore, if it had been intended to retain a general right on the part of the licensor to license his works in exactly the same exclusive domain as the licensee (notwithstanding the provisions of clauses 1(c) and 2(b)), then it is to be expected that the parties would have used much clearer words. Instead, a vague and ambiguous caveat has been inserted into a representation and warranty provision without any indication as to how that caveat is intended to operate in conjunction with the earlier provisions of the agreement or what the parties might have meant by “occasionally” or “legitimate creative purposes”.

723. Given the express purpose of this agreement as an exclusive licence agreement subject to the minor reservations in clause 2(b) (in respect of which Stability no longer takes issue), I consider that the only sensible and commercial reading of clause 9(b) is as a reference back to clause 2(b), as Getty Images submit. I accept of course that there is

no express reference back, and that has given me pause for thought. However, this reading at least ties the caveat in 9(b) to something tangible in terms of specific activity and gives it some meaning which is not entirely at odds with the remainder of the agreement. On this interpretation, any carve-out that clause 9(b) does make certainly cannot apply to the licensing of digital content and I remind myself that not all of a rightsholder's rights need to be the subject of an exclusive licence.

724. Even if I am wrong about that, it seems to me that clause 9(b) is so vague as to be unenforceable and that it cannot possibly affect the provisions of the rest of the agreement.
725. I reject Stability's case that Sample Licences #34-#38 are not on their face exclusive licences.

(iii) The Signature Point

726. Section 92(1) CDPA provides that an exclusive licence is "...a licence in writing signed by or on behalf of the copyright owner...". This raises a question as to the meaning of the word "signed" and in particular what is required in terms of content and form. Prior to trial, it appeared that there would be a substantial dispute over this issue, but by the time of trial, I did not understand the parties to disagree over the following propositions of law (adapted slightly from Getty Images' opening submissions):
 - i) The entity that signs must have sufficient authority to do so on behalf of the copyright owner;
 - ii) The licence must be "in writing" which the CDPA defines very broadly to cover just about any method or medium of fixation (see section 178 CDPA: "writing" includes any form of notation or code, whether by hand or otherwise, and regardless of the method by which, or medium in or on which, it is recorded..."). It follows that "signed" should be interpreted equally broadly because the licensor needs to be able to "sign" the document in the medium in which it exists.
 - iii) The principal function of a signature is to demonstrate an intention of the party to authenticate the document (see *Goodman v J Eban Ltd* [1954] 1 QB 550 per Evershed MR at 557 and *Bassano v Toft* [2014] EWHC 377 (QB), [2014] ECC 14, per Popplewell J at [42]).
 - iv) The requirement to "sign" in the CDPA must be interpreted as "always speaking" and should take account of technological developments which the legislators might not have foreseen, if they conform to the policy of the Act in question (see *News UK & Ireland Ltd v Commissioners for His Majesty's Revenue and Customs* [2023] UKSC 7; [2024] AC 89 at [28] and *Hudson v Hathway* [2022] EWCA Civ 1648, [2023] KB 345, per Lewison LJ at [56])).
 - v) Clicking on an "I Accept" tick-box on a website has been held to be sufficient to constitute a valid signature (see *Bassano v Toft* per Popplewell J at [42]-[44], a case involving agreements which had to be signed "in the prescribed form" under the Consumer Credit Act 1974).

vi) Similarly typing a forename or automatically generating an email signature in the footer of an email is sufficient to constitute signing a written document (see *Neocleouse v Rees* [2019] EWHC 2462 (Ch), [2020] P&CR 4 per Pearce HHJ at [55]-[57] and *Hudson v Hathway* per Lewison LJ at [67]).

727. Thus the parties did not disagree with the proposition (and I find) that typing “I Agree” or clicking a button labelled “I Accept” is sufficient to constitute signing of a licence within the meaning of section 92 CDPA. I understand this point of principle to apply to all licences falling within Groups K-N.

728. Further, by the time of trial, there were no specific issues in relation to the signing of the Contributor Agreements.

729. However, Stability contends that there is no evidence that Getty Images has any signatures for any of the iStock ASAs that pre-date 2012 – this applies specifically to #33 which dates back to 2006.

730. Stability accepts that after 2012, where a copyright owner has:

- i) typed “I agree” in a designated box beneath the terms of the exclusive iStock ASA described by Ms Malnar in her evidence; or
- ii) clicked a button labelled “I Accept” as described by Ms Varty in her evidence (as occurred in relation to around 100 contributors after February 2017),

that copyright owner will have signed the agreement. It is no longer contended that signing a licence agreement using an electronic form, whether by typing “I Agree” or by clicking on an “I Accept” button is not sufficient to satisfy the requirement that a licence be “signed” under section 92 CDPA. I also did not understand Stability to dispute that the use of “docusign” software (i.e. the use of an electronic signature), as occurred in respect of agreements entered into after 2017 (with the exception of those agreements which required the contributor to click a button labelled “I Accept” referred to above) was capable of amounting to a signature for the purposes of section 92 CDPA.

731. However, Stability contends that it is nonetheless necessary to consider each agreement on its own facts to determine whether the necessary acceptance in electronic form has happened or not.

732. In this context, Stability maintains a point on the absence of evidence of signatures in respect of SOCI Works A14 and A15, created respectively by Mr Devon Stephens and Mr Nikolay Pandev. In oral closing, Stability focussed only on A14, submitting that there was no evidence of any signature. In its written closing submissions, however, Stability continued to rely on the absence of evidence in respect of both A14 and A15 and so I deal with both in the following paragraphs of this section.

733. On the evidence of Ms Malnar, I find that from December 2012, when Ms Malnar was employed by the Fifth Claimant, the process for contributors uploading content to iStock and agreeing an exclusive licence required them to type “I agree” into the box underneath the terms of the exclusive iStock ASA. In her statement Ms Malnar said that she is aware through working at the Fifth Claimant that the process “did not change significantly from at least around 2010”, but she explained in cross examination that

this evidence was based on “the last update date” (i.e. the last revision date) for the iStock ASA. Ms Malnar accepted that if, during the period prior to her employment, there had been differences in the way in which the actual enrolment process worked, she would not have personal knowledge of that. In the circumstances, there is no detailed evidence as to the process for enrolment that was used by the Fifth Claimant prior to December 2012 and specifically no evidence as to how (if at all) any licences were signed prior to this date.

734. Furthermore, Ms Malnar did not know of any record of contributors submitting the appropriate form and typing “I agree” (at any time). She did say, however, (without any further explanation) that this was “programmed on the back end” – which I understood to mean that there would likely be some digital means of identifying whether individual contributors had provided the necessary confirmation of their agreement to the terms, although there is no other evidence to that effect in these proceedings. As far as I am aware, Getty Images has not disclosed any record of the individual contributors who have typed “I Agree” in order to confirm their acceptance of the iStock ASAs.

735. I turn then to look at the SOCI Works A14 and A15. Getty Images’ pleaded case is that both Mr Stephens and Mr Pandev are now subject to the latest iStock ASA from 2017 (iStock ASA #38). Their case (confirmed by the evidence of Ms Cameron and the Getty Images Editorial Workflow System) is that:

- i) Mr Stephens created SOCI Work A14 on or around 25 April 2007. According to Ms Cameron, it was uploaded to iStock on 16 November 2007. Mr Stephens entered into an iStock ASA with iStock International (a predecessor of the First Claimant) on 13 September 2006 on the terms of iStock ASA #33, but that from an unknown date in 2017 he has been subject to the terms of iStock ASA #38. No signed agreement has been disclosed from 2006, 2017 or any other date.
- ii) Mr Pandev created SOCI Work A15 on 17 October 2019. He entered into an iStock ASA on the terms of #37 on 26 February 2016 and that from an unknown date in 2017 he has been subject to the terms of iStock ASA #38. No signed agreements have been disclosed.

736. In both cases, hearsay evidence is available in the form of written letters provided by Mr Stephens and Mr Pandev at the request of Ms Cameron. In his letter dated 19 March 2025, Mr Stephens says that he believes he signed up for iStock in 2006 and that:

“If I recall correctly there was an application process that included uploading a few images that were checked for technical quality before I was approved to upload more. **Terms would have been accepted through the website**, the image was uploaded through the website, along with model releases. Originally iStock was not exclusive, but **I joined iStock’s exclusivity program before I uploaded [A14]**”.

737. Stability did not suggest that I should attach no weight to this evidence and accordingly, I find that, on balance, the terms on which SOCI Work A14 were supplied by Mr Stephens to iStock International (#33) were signed by Mr Stephens using some form of digital acceptance. There is no evidence from Mr Stephens as to whether he has agreed

to the terms of iStock ASA #38 and no evidence as to whether he was even contributing content as at the date of this agreement.

738. In a letter from Mr Pandev dated 20 March 2025, he says this:

“I became a contributor to iStock (which I understand is part of Getty Images) in around January 2015 and I became an exclusive contributor in around 2016 (before I uploaded [A15]). As far as I can recall, the process of becoming an exclusive contributor was very straightforward and involved me agreeing to the exclusive contributor terms online via iStock’s website”.

739. Again, it appears to me to be reasonable to find (based on this evidence) that Mr Pandev agreed to the iStock terms in 2016 using some form of digital acceptance. However, A15 was not taken until 2019, after Mr Pandev appears to have been moved to a new agreement. There is nothing in his letter which refers to this later agreement and thus no evidence that he accepted the terms of that agreement.

740. Ms Malnar’s evidence, which I accept, is that between 2015 and 2017 Getty Images engaged in a project to bring iStock and Getty Images contributors onto the same system. This appears to be the rationale behind migrating iStock contributors onto the new iStock ASA #38. Ms Malnar’s evidence under cross examination was that contributors could not upload images until they had agreed to the new terms and that they would do this by typing “I Agree” and then submitting the form. However, the submission of the form could be deferred or ignored and after a particular date contributors who had not yet confirmed their agreement would be deemed to have accepted and continue to upload content. This would happen regardless of the fact that those contributors had not formally “signed” the new terms by submitting the online form. In closing, Getty Images accepted that in such circumstances the contributor/licensor will not have signed the new licence agreement onto which he or she is being migrated within the meaning of section 92 CDPA, albeit that they contended that “the overwhelming likelihood is that, even where the licensor deferred, the licensor would then accept the terms by typing “I Agree”.

741. On balance, in light of Ms Malnar’s new evidence only given during cross examination, I accept that there is insufficient available evidence on which I can find that Mr Stephens and Mr Pandev in fact “signed” the new terms in 2017 by typing “I Agree” and submitting the form. Without any evidence of what course they in fact took, and without knowing whether (in general) the vast majority of contributors signed the new terms (as opposed to merely leaving them and then being deemed to have accepted them), I reject the suggestion that I can find that the “overwhelming likelihood” is that they expressly accepted those terms. Importantly, as it seems to me, there is nothing in their respective letters to that effect.

742. In any event, as I have said, I accept that there is evidence of Mr Stephens having “signed” iStock ASA #33 and of Mr Pandev having “signed” iStock Agreement #37 for the purposes of the CDPA.

743. Finally, I agree with Stability that the question of whether licences have been “signed” is fact sensitive and can only be determined by reference to the individual licences. Certainly there is no basis at this liability trial on which I can determine how many

licence agreements are in fact exclusive because they have been “signed”. Although the Order of 17 January 2025 provided that the liability trial should determine “whether the Schedule 1 Agreements are exclusive licences within the meaning of s.92 CDPA (the Licensing Issue”) by reference to a sample”, now that I have heard the available evidence, I do not consider that this evidential point as to signature (as opposed to legal points about the true interpretation of the agreements) can properly or fairly be determined in this way. Specifically, the fact that I have found that a sample licence has been signed by an individual contributor, does not mean that one can fairly extrapolate that finding across all other agreements falling within the same group of licences which will involve numerous other contributors, and I decline to do so. I did not understand Getty Images to disagree with this approach.

Amendments and variations to Licence Agreements

744. A point is taken by Stability as to clauses in the Sample Licences permitting amendments to agreements. Stability accept that from June 2014 all Contributor Agreements and all iStock ASAs contain a clause permitting Getty Images or others unilaterally to amend these agreements – something which may be done through posts or updates on the Getty Images Websites. However, Stability says that prior to that date all agreements (namely #1-13 and 30-32) can only be varied by the written agreement of the parties (see by way of example clause 4.1 of #11).
745. The short answer to this point is that it has not been adequately pleaded (a reference at paragraph 15AB(iii) in the Defence in parenthesis is not sufficient in my judgment). Getty Images also point out that each Contributor Agreement expressly supersedes and cancels all previous agreements between the parties in any event (which means that the date on which clause 4.1 was introduced is of no significance). I did not understand Stability to pursue any remaining point of substance in closing. Its submissions on variation and amendment do not in any event appear to affect any of the exclusively licensed works in issue in this trial.

The Sixth Claimant’s Licence Agreements

746. For completeness I observe that each of the relevant licences entered into by the Sixth Claimant is a Sample Licence within one of the relevant groups and accordingly I need make no separate determinations on the question of whether the Sixth Claimant’s licence agreements are exclusive. In any event, as Getty Images point out, the Sixth Claimant is a claimant in these proceedings in its own right and so this issue is, to all intents and purposes, purely academic.

(K) REMAINING OUTSTANDING ISSUES

Number of works used in training

747. One of the outstanding issues for trial is the number of Visual Assets and Copyright Works used in training. However, there is a dispute between the parties as to whether this is an issue that should still be determined, notwithstanding that Getty Images have abandoned the Database Rights Claim, the Training and Development Claim and the Outputs Claim.

748. Getty Images submit that the numbers are relevant to at least three issues: (i) secondary infringement of copyright “because Stability’s admissions do not cover v1.1-1.3”; (ii) trade mark infringement and passing off “because the likelihood of a watermark appearing in an output is dependent on the frequency with which the watermark appears in the training data”; and (iii) additional damages “because when the court comes to consider the issue of flagrancy, a relevant factor to weigh in the balance will be the scale of any infringement”.

749. Stability accepted at the PTR that “many” Copyright Works were used to develop and train Stable Diffusion. In its statements of case it has admitted that “at least some” Copyright Works were used in training v2.0, that at least some images from the Getty Images Websites were used during training; and that one or more of Sample Works A1-A17 were used in the training of v1.4, v2.x, all versions of SD XL and v1.6. It has also admitted for the purposes of these proceedings that there are approximately 12.3 million Visual Assets in LAION 2B-en. It contends that there is no reason for the court to try to make any specific findings going beyond these admissions as to either Visual Assets or Copyright Works. I am inclined to agree. As to the former, it is in any event unclear why a finding as to Visual Assets could continue to be of relevance given the abandonment of the Database Rights Claim. As to the latter, Stability correctly points out that any attempt to make findings on the number of Copyright Works would raise complex issues - not least because Getty Images’ case requires them to satisfy the court that the assets used in training were images whose copyright was owned by one of the Claimant companies or whose copyright has been licensed to one of the Claimant companies under one of the Annex 3 Licence Agreements. Insofar as Getty Images is inviting the court to determine the rights ownership position for each work, that is a fact sensitive question for each work and plainly impractical.

750. I decline to make any finding as to the number of Visual Assets or Copyright Works used in training. In addition to the points made by Stability and set out above, and taking each of Getty Images’ reasons for inviting me to do so in turn:

- i) While it is true that Stability’s admissions do not cover v1.1-1.3, I cannot see the relevance of a determination as to the number of Copyright Works in connection with any aspect of the Secondary Infringement Claim, which has, in any event, been tried by reference to identified Sample Works and the SOCI Works only. I have not needed to address the question of the number of Copyright Works used in training in connection with my determination of that claim, which I have dismissed on the point of construction of the CDPA.
- ii) It may be the case that the likelihood of a watermark* appearing on a synthetic image output is dependent upon the frequency with which the watermark appears in the training data. However, Getty Images (i) have not sought to advance any probabilistic case in these proceedings as to the likelihood of watermarks* appearing; and (ii) have no evidence as to the correlation between the images bearing watermarks being used in training and the production of synthetic images bearing watermarks*. Furthermore, although it is clear that filters were used during the development of various Models, including as a means of filtering watermarks from training data, there is no evidence as to how effective these may have been. That a particular number of Copyright Works may have been in the original training dataset does not mean that the Model was in fact trained on these images once any filters were applied. This was confirmed

by Mr Wagrez in his evidence and the use of filters (at least in relation to the later Models) is clear from some of the contemporaneous documents to which I have already referred.

- iii) For reasons to which I shall return in a moment, this is not a case in which I consider an award of flagrancy damages to be appropriate.
- 751. Finally, and in addition to the fact that I cannot see the need to make any finding as to the number of Visual Assets or Copyright Works in the training data, I do not consider that the available evidence would enable me to do so. I note, in particular that I could attach little weight to Ms Cameron's fourth witness statement owing to the way in which that evidence was produced and Ms Cameron's limited understanding of the steps taken to put it together. In their written submissions, Getty Images put their case on numbers no higher than that "it cannot sensibly be suggested that the development and training of each of the sub-versions of Stable Diffusion v1.x, v2.x and XL involved substantially fewer than millions of Visual Assets and Copyright Works". I can see no reason to make such a vague (and ultimately irrelevant) finding in the circumstances of this case. Given the way in which the case has been advanced, I cannot see that it would assist me in arriving at any reliable conclusion as to the scale of watermark* generation on synthetic images by users in the UK.
- 752. For all the reasons I have given, I decline to make any such finding.

Additional Damages

- 753. There is no dispute between the parties as to the relevant legal principles. Getty Images referred me to the Intellectual Property (Enforcement, etc.) Regulations 2006 which implement Directive 2004/48/EC and to section 97(2) CDPA (which is concerned only with copyright infringement).
- 754. Getty Images' pleaded case in support of additional damages relies upon three matters: first that Stability knows or has reason to believe that Stable Diffusion is an infringing copy of the Copyright Works; second that Stability has acted with a "cavalier attitude" to the rights of the Claimants, "infringing numerous different intellectual property rights on a blatant and widespread scale"; and third that Stability has relied upon the infringement of copyright and database right for the successful operation of its business and Stable Diffusion has been, or has the potential to be, used by end users who would otherwise have used the Getty Images Websites.
- 755. Given that I have dismissed the Secondary Infringement Claim and that both the Outputs Claim and the Training and Development Claims were abandoned at trial, I cannot see how either the first or third of the particulars pleaded by Getty Images goes anywhere. As Stability submits, the position is that the Models are producing novel content in the United Kingdom against which Getty Images have no claim. Where the court cannot conclude that the training was an infringement (because it took place in the United States, or elsewhere) and has found that there have been no acts of secondary infringement, there can be no basis for additional damages for infringement of copyright under section 97(2) CDPA. I dismiss the claim for additional damages under that section.

756. As for the second of Getty Images' particulars, it cannot possibly be maintained in light of my findings that Stability has infringed "numerous different intellectual property rights on a blatant and widespread scale". While I have found instances of trademark infringement, I have been unable to determine that these were widespread, or that they continued beyond the release of v2.x. In the circumstances, there is, in my judgment, no basis whatsoever for a claim for additional damages under the broader provisions of the 2006 Regulations. Getty Images may be able to maintain such a case in the jurisdiction where the Model was in fact trained, but there is no basis for that case in this jurisdiction.

(L) CONCLUSION

757. In summary, although Getty Images succeed (in part) in their Trade Mark Infringement Claim, my findings are both historic and extremely limited in scope. The Secondary Infringement Claim fails.

758. In a little more detail, my findings on the key issues are as follows:

- i) Stability bears no direct liability for any tortious acts alleged in these proceedings arising by reason of the release of v1.x Models via the CompVis GitHub and CompVis Hugging Face pages.
- ii) the question of trade mark infringement arises only:
 - a) in respect of the generation of Getty Images watermarks* and iStock watermarks* by v1.x Models (in so far as they were accessed via DreamStudio and/or the Developer Platform);
 - b) in respect of the generation of Getty Images watermarks* by v2.x Models.
- iii) There is no evidence of a single user in the UK generating either Getty Images or iStock watermarks* using SD XL and v1.6 Models. Thus no question of trade mark infringement arises in respect of these Models and that claim, in so far as it relates to them, is dismissed.
- iv) As to Getty Images' claim under section 10(1) TMA:
 - a) Getty Images succeed in respect of iStock watermarks* generated by users of v1.x (in so far as the Models were accessed via DreamStudio and/or the Developer Platform) in that infringement of the ISTOCK Marks pursuant to section 10(1) TMA has been established. This success is however based specifically on the example watermarks* shown on the Dreaming Image and the Spaceships Image – the latter having been generated by Model v1.2. Given the way in which the case has been advanced, it is impossible to know how many (or even on what scale) watermarks* have been generated in real life that would fall into a similar category.

- b) Getty Images fail in respect of Getty Images watermarks*, there being no evidence of infringement of the Getty Images Marks under section 10(1) TMA. That claim is dismissed.
- v) As to Getty Images' claim under section 10(2) TMA:
 - a) Getty Images succeed in respect of iStock watermarks* generated by users of v1.x (in so far as the Models were accessed via DreamStudio and/or the Developer Platform) in that infringement of the ISTOCK Marks pursuant to section 10(2) TMA has been established. This success is based specifically on the example watermarks* shown on the Dreaming Image and the Spaceships Image - the latter having been generated by Model v1.2.
 - b) Getty Images succeed in respect of Getty Images watermarks* generated by users of v2.x in that infringement of the Getty Images Marks pursuant to section 10(2) TMA has been established. This success is based specifically on the example watermark* on the First Japanese Temple Garden Image, generated by Model v2.1.
- Again, it is impossible to know how many (or even on what scale) watermarks* have been generated in real life that would fall into a similar category.
- vi) Getty Images' claim under section 10(3) TMA is dismissed.
- vii) For reasons I have explained, I have declined to address Getty Images' allegation of passing off.
- viii) Getty Images' claim of secondary infringement of copyright is dismissed. Although an "article" may be an intangible object for the purposes of the CDPA, an AI model such as Stable Diffusion which does not store or reproduce any Copyright Works (and has never done so) is not an "infringing copy" such that there is no infringement under sections 22 and 23 CDPA.
- ix) As to other remaining issues:
 - a) On copyright subsistence/ownership, Getty Images have failed to make out title to the copyright in SOCI Works A3 and A4, but have established their title to the copyright in SOCI Works A9, A10 and A11.
 - b) On the Licensing Issue:
 - i) Sample Licences #2, 3, 10, 11, 13, 30 and 32 are not exclusive licences under section 92 CDPA;
 - ii) Sample Licences #17, #19, #34-38 are exclusive licences under section 92 CDPA.
 - c) I make no finding as to the number of Visual Assets or Copyright Works used in training Stable Diffusion.

d) This is not a case in which I consider an award of additional damages to be justified.

759. Any consequential matters arising in light of this judgment will be dealt with at a hearing in due course.

Appendix A

Glossary of Terms

AI	Artificial Intelligence.
ADD	Adversarial Diffusion Distillation. A combination of aspects of Generative Adversarial Networks (GANs) and Diffusion Models.
API	Application Programming Interface. A defined protocol that allows two computer programs to communicate and exchange data.
Autoencoder	An autoencoder is a system that can encode an image from a “pixel space” representation of an image into a representation in “latent space” (the encoder role) or from latent space back into pixel space (the decoder role).
AWS Cluster	Amazon Web Services Inc cloud hosting and processing resources.
Batch/Batch size	In the context of training Stable Diffusion, the number of training images processed by the model during a single training step or iteration.
Bias	See also “weights”. A constant added to the sum of the weighted inputs of a neuron (or node) in a neural network which affects the output. The weights and biases are sometimes just called the “weights” of a model.
Checkpoint	A file recording the current state of training of a machine-learning model. Further training can continue from that checkpoint. Versions of Stable Diffusion may be released at a particular “checkpoint” and then training continued for later versions. Versions v1.1, v1.2, v1.3 and v1.4 of Stable Diffusion are both checkpoints and versions.
CLIP/CLIP text encoder	Contrastive Language-Image Pre-Training. It is a deep-learning model that has learned the relationship between images and text. It has an image encoder and a text encoder.
Diffusion Model	An AI system that generates an output, such as an image, by a process of starting with an image consisting of noise, or an image to which noise has been added, and progressively removing the noise to generate an output image.
Deep Learning	A procedure to optimise both the decision function and the underlying input representation based on data.
Denoising	The process of removing noise from an image.

Epoch	A complete pass of all the data in a training dataset during a training process.
GANs	Generative Adversarial Networks. Unlike Diffusion Models, GANs can produce images in a single step but include drawbacks such as lower sample quality.
Generalisation	A network such as Stable Diffusion must be able to perform well on new data. This means it must generalise – i.e. recognise and understand the general patterns and rules in the training data and be able to apply them in a different context.
Generative AI	Machine Learning models that generate content such as images or text as their output.
GPU	Graphics Processing Unit.
Gradient Descent	A method of training in which a model plots a route from a given starting point (i.e., the model's unoptimized starting loss function) to the lowest point (i.e. a minimum of the cost function) by taking the steepest descending local gradient.
Inference	The process of running a machine learning model to produce outputs (images in the case of Stable Diffusion).
Iteration	The process by which the model being trained generates an output from a training data sample (or batch of samples), calculates the loss function, calculates the gradient of the loss function on a layer-by-layer basis, and updates the weights to achieve the gradient descent step.
Latent space	An image in “latent space” is an image that has been converted from an image in pixel space using an autoencoder. Latent space is a compressed, representative form of the pixel space image that is more memory- and computationally-efficient.
Layers	A collection of neurons or nodes. In a neural network, information flows from layer to layer. Each layer is composed of a collection of neurons or nodes. There is an “input layer” and an “output layer”, and there may be intermediate layers, known as “hidden layers”.
Loss Function	A quantification of the difference between the target or expected output of a model and the actual output. Loss function is used in a model's training process in determining adjustments to its weights that should result in a reduced loss function.

Machine Learning	Given a task to solve, a machine-learning system is designed to infer a solution from data without being instructed explicitly on how to solve the problem. The system is trained by providing it with data from which it can “learn” to carry out tasks not explicitly provided for by the software engineer’s design. Machine learning is part of the wider field of AI.
Memorisation	This process occurs in image models such as Stable Diffusion when the model is over-exposed to one specific training image linked to a particular concept. The process of memorisation undermines the model’s ability to generalize and generate new content when prompted with a text prompt for that concept.
Model Card	A short, standardised document that provides detailed information about the Model.
Neural Network	A machine learning model consisting of a network of neurons (or nodes).
Noise	In the context of an image, noise can take the form of random pixel values. An existing image can be partially corrupted with noise, or an image can consist entirely of noise.
Open source	Software whose source code is made freely available for anyone to use.
Overfitting	A situation where the model produces outputs that fit too closely to the training samples rather than generalising from them.
Pixel/Pixel space	A digital image suitable for human viewing may be represented by a two-dimensional array of “pixels”. An image in “pixel space” is an image in this form.
Prompt	A descriptive text caption or an image (or both) input by a user to condition the output of Stable Diffusion.
Sample/Sampling	The value drawn from the probability distribution. Drawing a value from a specified probability distribution.
Stochastic system	A system in which the same input may not produce the same output because of some randomness or probabilistic behaviour of the system.
Seed/Seed value	In Stable Diffusion a “seed” is a latent image used as the starting point for the inference process. The seed may be pure random noise or a user image to which noise has been added. In the case of text-to-image inference, the user can specify the “seed value” (a numeric value) in which case the same

	seed image will be used each time, and the output will be identical or practically so.
Step	In the context of training a model like Stable Diffusion a “step” is the process of adjusting the model weights at the end of each iteration. It can also be used synonymously with “iteration”. See also “Iteration”.
Text Encoder	The text encoder takes a sequence of tokens (numbers representing words or parts of words) and converts them using a deep learning model to form a “text-embedding” or an “embedding” which is designed to represent the semantic meaning of the words.
Training	The process by which a machine-learning model “learns” its model weights from user-provided data.
Tokens/Tokeniser	A tokeniser converts text into tokens, which are numbers representing words or parts of words.
U-Net	A neural network that, in the case of Stable Diffusion, is responsible for “denoising” the image in latent space. The denoising process may be conditioned by user prompts to influence the resulting image.
URL	Uniform Resource Locator – a hyperlink (via a web address) to the original internet location at which the referenced image is made available to the public by the website operator.
Vector File	A resolution independent depiction of a visual image that can be scaled up or down and is often used for illustrations.
VRAM	Video Random Access Memory.
Weights	In a neural network a neuron (or node) has inputs (which may come from outside the network or from a previous layer of the network) and an output. A “weight” determines the strength of the effect a particular input has on the output. See also “bias”. The weights and biases are sometimes just called the “weights” of a model. The aim of training is to adapt the weights and biases such that the model produces desired outputs.

Appendix B

The relevant context in which the average consumer will encounter the watermarks*

1. I consider in this Appendix the manner in which consumers are, or would be, exposed to the Signs in issue. I need only undertake this exercise in relation to v1.x and 2.x (owing to factual findings made in the Judgment ruling out trade mark infringement in respect of SD XL and v1.6). Ultimately I have determined that none of the matters relied upon by Stability in respect of the means of access as relevant context negates the scope for infringement.
2. I shall address the following scenarios:
 - (i) a consumer generating content through a locally downloaded copy of Stable Diffusion v2.0;
 - (ii) a consumer generating content through a DreamStudio account on Stable Diffusion v1.4, and v2.0; and
 - (iii) a consumer generating content through an API available on the Developer Platform in relation to Stable Diffusion v1.x or v2.x.
3. In so doing, I adopt much of what Stability says in its written submissions about the facts (which I did not understand to be in dispute).

Local Downloads via GitHub and Hugging Face

4. I have dealt in the main Judgment with the potential for this to take place in two ways: using either the Direct Download Method or the Diffusers Method.
5. The notional assessment of the court requires it to consider the average consumer generating images through each of these mechanisms, as well as what the user understands by way of context from the GitHub and Hugging Face pages themselves, which are not identical for the Models in issue.

Stable Diffusion v2.x

6. The Stability GitHub page for v2.x includes the following features:

- (i) The requirements for downloading the Model are identified as either (a) an update to an existing latent diffusion environment, with instructions for running the appropriate code; or (b) for “more efficiency and speed” installation of the xformers library with guidance on installation.
- (ii) A General Disclaimer in the following terms:

“Stable Diffusion models are general text-to-image diffusion models and therefore mirror biases and (mis-) conceptions that are present in their training data. Although efforts were made to reduce the inclusion of explicit pornographic material, we do not recommend using the provided weights for services or products without additional safety mechanisms and considerations. The weights are research artifacts and should be treated as such. Details on the training procedure and data, as well as the intended use of the model can be found in the corresponding model card. The weights are available via the StabilityAI organization at Hugging face under the CreativeML Open rail +-M Licence”.

- 7. The CreativeML Open Responsible AI M Licence (in which copyright is identified as belonging to “Stability AI and contributors”) contains the following provisions:

“Multimodal generative models are being widely adopted and used, and have the potential to transform the way artists, among other individuals, conceive and benefit from AI or ML technologies as a tool for content creation.

Notwithstanding the current and potential benefits that these artifacts can bring to society at large, there are also concerns about potential misuses of them, either due to their technical limitations or ethical considerations...

This license governs the use of the model (and its derivatives) and is informed by the model card associated with the model.

NOW THEREFORE You and Lessor agree as follows:....

5. Use-based restrictions. The restrictions set forth in Attachment A are considered Use-based restrictions. Therefore you cannot use the Model and the Derivatives of the Model for the specified restricted uses. You may use the Model subject to this License, including only for lawful purposes and in accordance with the License. Use may involve creating any content with

finetuning, updating, running, training, evaluating and/or reparametrizing the Model...

6. The Output You Generate. Except as set forth herein, Licensor claims no rights in the Output You Generate using the Model. **You are accountable for the Output you generate and its subsequent uses.** No use of the Output can contravene any provision as stated in the License" (**emphasis added**).
8. The "Use-based Restrictions" in Annex A are stated as follows:

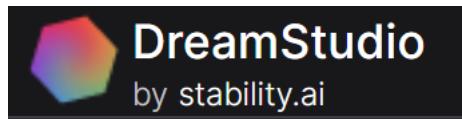
"You agree not to use the Model or Derivatives of the Model:

 - In any way that violates any applicable national, federal, state, local or international law or regulation;
 - For the purpose of exploiting, harming or attempting to exploit or harm minors in any way;
 - To generate or disseminate verifiably false information and/or content with the purpose of harming others;
 - To generate or disseminate personal identifiable information that can be used to harm an individual;
 - To defame, disparage or otherwise harass others...”.
9. The Stability Hugging Face page for each of the checkpoints of v2.x includes Model Cards which have the following features:
 - (i) The model license is again stated to be subject to a CreativeML Open RAIL++-M License.
 - (ii) Notes recommending the installation of xformers for better performance together with a particular upgrade “if you have a low GPU RAM available”.
 - (iii) Information about the uses of the Model, its training data and misuses, malicious use and out of scope use. In particular it includes the further additional features:
 - a. It is explained that the Model is “intended for research purposes only” and possible research areas and tasks are then identified, namely (a) safe deployment of models which have the potential to generate harmful content; (b) probing and understanding the limitations and biases of generative models; (c) generation of artworks and use in design and other artistic processes; (d) applications in educational or creative tools; (e) research on generative models.

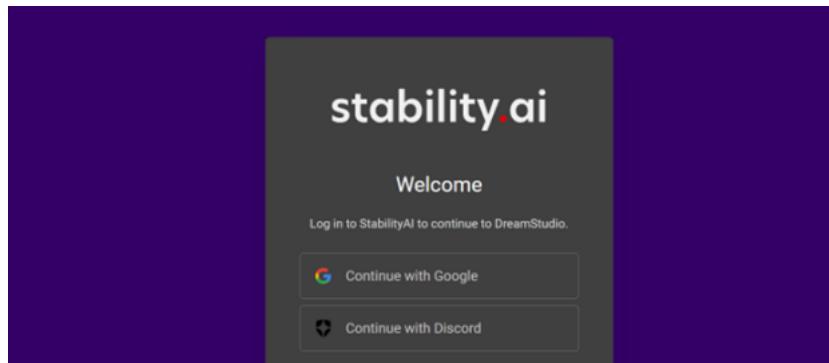
- b. Under the heading “Misuse and Malicious Use” it is explained that “[t]he model should not be used to intentionally create or disseminate images that create hostile or alienating environments for people. This includes generating images that people would foreseeably find disturbing, distressing or offensive; or content that propagates historical or current stereotypes”. Misuse is said to include (but is not limited to): (a) generating demeaning, dehumanizing, or otherwise harmful representations of people or their environments, cultures, religions, etc; (b) intentionally promoting or propagating discriminatory content or harmful stereotypes; (c) impersonating individuals without their consent; (d) sexual content without consent of the people who might see it (e) mis- and disinformation; (f) sharing copyrighted or licensed material in violation of its terms of use; and (g) sharing content that is an alteration of copyrighted material or licensed material in violation of its terms of use.
- c. Under the heading “Out of Scope Use” it is expressly stated that “[t]he model was not trained to be factual or true representations of people or events, and therefore using the model to generate such content is out of scope for the abilities of this model”.
- d. Under the heading “Limitations” it is stated (amongst other things) that “the Model does not achieve perfect photorealism” and that “the model was trained on a subset of a large scale dataset LAION-5B which contains adult, violent and sexual content. To partially mitigate this, we have filtered the dataset using LAION’s NSFW detector”.

DreamStudio (v.1.4 and 2.0)

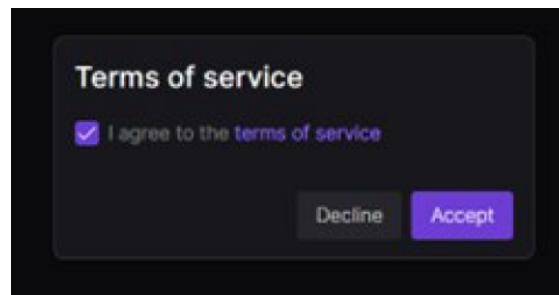
- 10. Stability addressed the details of accessing the Model via DreamStudio in its skeleton arguments by reference to v1.6. However, I did not understand there to be any significant difference between the use of DreamStudio as an access mechanism for users of v1.x, v2.x and v1.6. None was identified.
- 11. The web address for DreamStudio is beta.dreamstudio.ai. The ‘favicon’ is a hexagonal colour palette. The hexagonal colour palette is again repeated as a page logo in the top left of the screen next to a title ‘DreamStudio’ under which it reads ‘by stability.ai’.



12. In order to access the service, the average consumer/user needs to create an account and log in to that account. It requires an email sign-up and confirmation of a verification email sent by Stability to the user's email account.
13. Logging into the account, the user is again faced with a large stability.ai logo:



14. In order to proceed, the user has to agree to the Terms of Service:



15. By this point it is clear to the user that they are interacting with a service provided by Stability.ai, although I accept Getty Images' submission that there is no evidence (or pleading) that Stability has any goodwill or reputation in the names "DreamStudio" or "Stability ai" or in the favicon.

16. The terms of service include the following statement:

"The Services are provided by Stability AI Ltd... to provide creative generative tools for our users to express their creativity to produce user-generated images from text prompts ("Content"). The Services use an artificial intelligence ("AI") tool to generate this Content. This use of AI is relatively new and still evolving. As a result, while we have taken — and continue to take — efforts to preclude your creation of extreme content, we cannot guarantee the suitability or appropriateness of the resulting images you generate. You are solely responsible for your use of the

Services, including your text prompts, generation of Content, and the consequences of your Content Sharing (as defined below).

...

Your use of the Services is subject to the terms, conditions, and use restrictions in the Creative ML OpenRAIL-M license (available at <https://huggingface.co/spaces/CompVis/stable-diffusion-license>) (the “License”). You must review and follow the terms of the License.

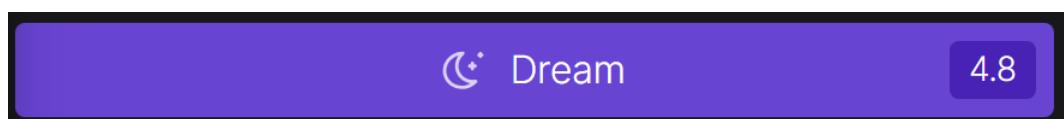
...

Accordingly, you may not, or assist any other person to:

- Generate any Content or engage in Content Sharing that a reasonable person could find obscene, lewd, lascivious, offensive, pornographic, indecent,...
- Generate Content or engage in any Content Sharing that is intentionally misleading, false, or otherwise inappropriate or with the purpose of harming others, regardless of whether the Content or its dissemination is unlawful”.

17. The License terms are in materially the same form as those referred to above. They inform the user that the outputs of the model are their own responsibility and that they may contain content which is false, misleading etc.

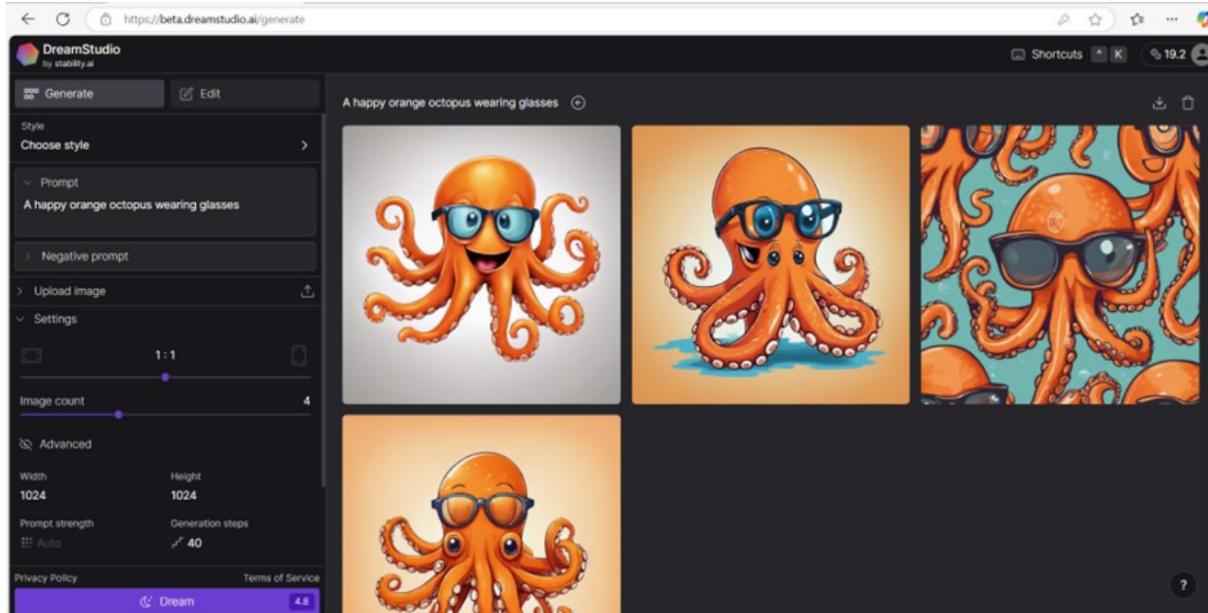
18. The user is required to purchase credits to allow them to use the model. They then select a model and input their prompt. The process of inference is started by clicking the button marked ‘Dream’. Within the button is embedded a counter indicating the number of credits that will be used.



19. The user is given the option of how many images to generate at once, the default value being 4. Various other parameters are displayed, such as the number of generation steps.

20. A delay occurs before the resulting images are provided, in which the hexagonal colour palette logo is displayed within the pane which will become a synthetic

image. The (default) four synthetic images are then presented together in the web interface:



The Developer Platform/API (v1.x and v2.x)

21. The method of accessing the Developer Platform (platform.stability.ai) combines features of the DreamStudio and GitHub access mechanisms, in that it requires technical expertise to access the API through the appropriate code (as illustrated in detail in exhibit AKK-5 to Ms Kabba-Kamara's statement), but also requires the user to access the Stability website and create an account with Stability.
22. When accessing the Stability website, the user will see the same welcome message as appears for a user wishing to use DreamStudio and, having clicked "login", he or she can then choose to create an account (which is verified by an email from platform.stability.ai) in order to continue to the Developer Platform. Upon accepting that the Developer Platform can access the account, a user is presented with a pop-up click option inviting acceptance of the Stability AI API Terms of Service. These are said to "govern your access and use of these APIs" and provide that:

"The Services are provided by Stability AI Ltd...to provide creative generative tools for users to express their creativity to produce user-generated images from text prompts ("Content"). The Services use an artificial intelligence ("AI") tool to

generate this Content. This use of AI is relatively new and still evolving. As a result, while we have taken – and continue to take – efforts to preclude your creation of extreme content, we cannot guarantee the suitability or appropriateness of the resulting images you generate. You are solely responsible for your use of the Services, including your text prompts, generation of Content, and the consequences of your Content Sharing...”

23. The API Terms of Service available in the evidence are expressly said to have been “last updated December 14 2022”. I can only infer that there are no material changes in the terms of service since the date of launch of Stability’s API offering on 25 November 2022.
24. The API Terms of Service state that the Services are provided based on models licensed under the CreativeML Open Rail +-M License and a link is provided. As with the reference to this License on the GitHub and Hugging Face pages there is no express requirement that a user read or accept the terms of the License.