

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 Licensor 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.